# Classifying Interplanetary Discontinuities Using Supervised Machine Learning

**Daniel Dumitru[1] and Costel Munteanu[2]**

[1]National Institute for Laser, Plasma and Radiation Physics, Măgurele, 077125, Romania.

[2]Institute of Space Science, Măgurele, 077125, Romania.

Corresponding author: Costel Munteanu (costelm@spacescience.ro).

## Key Points:

- In-situ magnetic field observations from Cluster 1 spacecraft are used to compile a database of events
- Localized time-scale images are created for each event, and supervised machine learning is used to classify them
- An investigation of average occurrence rates versus spacecraft location, demonstrates the validity of our results

## Abstract

Directional discontinuities (DDs) are defined as abrupt changes of the magnetic field orientation. We use observations from ESA's Cluster mission to compile a database of events: 4216 events are identified in January-April 2007, and 5194 in January-April 2008. Localized time-scale images depicting angular changes are created for each event, and a preliminary classification algorithm is designed to distinguish between: simple - isolated events, and complex - multiple overlapping events. In 2007, 1806 events are pre-classified as simple, and 2410 as complex; in 2008, 1997 events are simple, and 3197 are complex. A supervised machine learning approach is used to recognize and predict these events. Two models are trained: one for 2007, which is used to predict the results in 2008, and vice-versa for 2008. To validate our results, we investigate the discontinuity occurrence rate as a function of spacecraft location. When the spacecraft is in the solar wind, we find an occurrence rate of ~2 DDs per hour and a 50/50 % ratio of simple/complex events. When the spacecraft is in the Earth's magnetosheath, we find that the total occurrence rate remains around 2 DDs/h, but the ratio of simple/complex events changes to ~25/75 %. This implies that about half of the simple events observed in the solar wind are classified as complex when observed in the magnetosheath. This demonstrates that our classification scheme can provide meaningful insights, and thus be relevant for future studies on interplanetary discontinuities.

## 1. Introduction

Abrupt changes in the orientation of the interplanetary magnetic field (IMF), referred to as directional discontinuities (DDs), are ubiquitous structures in the solar wind. With an average occurrence rate at Earth of about two DDs per hour (e.g., Newman et al., 2020), these structures represent an omnipresent source of variability for the interplanetary plasma environment. DDs are known to trigger geomagnetic storms and magnetospheric substorms, with significant impact on ground-based and spaceborne technologies (e.g., Tsurutani et al., 2011). They can be used, for example, to estimate the solar wind propagation time from an upstream solar wind monitor to a downstream target (e.g., Mailyan et al., 2008; Haaland et al., 2010; Munteanu et al., 2013).

The term "directional discontinuity" was originally introduced by Burlaga (1969) to denote a variation of IMF direction larger than 30 degrees in less than 30 seconds. Many previous studies used the limit of 30° to distinguish between the population of turbulent fluctuations (characterized by directional changes below the limit) and the population of discontinuities (above the limit; see, e.g., Borovsky et al., 2008). This definition was the starting point for multiple detection algorithms. Li (2008), for example, describe a rather complex algorithm to identify discontinuities based on directional changes. Borovsky (2010) used a similar approach to identify solar wind DDs, and then studied their effects on the power spectrum. Chian & Muñoz (2011) used the Li (2008) detection method, and investigated the relation between discontinuities, turbulence, and magnetic reconnection at the leading edge of an interplanetary coronal mass ejection. The detection algorithm of Li (2008) was further developed by Miao et al. (2011), who introduced a way of automatically estimating the discontinuity thickness.

There are other ways of identifying magnetic field discontinuities. Vasquez et al. (2007), for example, developed a detection algorithm which is independent of directional changes, and instead relies on changes of the amplitude of magnetic field components. They used their algorithm to identify a large number of events, and found that the occurrence rate of solar wind discontinuities from their algorithm is comparable with that from algorithms based on directional changes. Tsurutani and Smith (1979) were among the first to develop a detection method based on changes of the amplitude of field components, and showed that it provides similar results to directional change-based methods. Burkholder and Otto (2019) introduced yet another detection algorithm based on amplitude changes. A notable contribution is the method called partial variance of increments (PVI; Greco et al., 2008; Greco & Perry, 2014; Greco et al., 2016; Greco et al., 2018). Greco et al. (2008) compared the results from PVI with those obtained using the Tsurutani and Smith (1979) method, and found that the two sets of results are remarkably similar, suggesting that most of the events identified by the two methods are the same.

Due to various computational difficulties encountered when implementing automated detection algorithms, even recent studies still use visual inspection to identify discontinuities (Mailyan et al., 2008; Munteanu et al., 2013; Artemyev et al., 2018, 2019a, 2019b). Note that even the (partially) automated detection algorithm of Burkholder and Otto (2019) still uses visual inspection to eliminate events that are not clearly isolated from other structures in the time series. For relatively small datasets, detection by visual inspection can be acceptable, but, for large-scale studies, visual inspection is certainly not suitable.

Magnetohydrodynamics defines two idealized classes of discontinuities: (a) stationary structures, i.e. discontinuities that do not propagate with respect to the ambient plasma (tangential (TDs) and contact discontinuities), and (b) propagating discontinuities (rotational discontinuities (RDs) and shocks). The most frequent small-scale discontinuities in

interplanetary space are the abrupt changes in the direction of the magnetic field, predominantly expected for TDs and RDs (e.g., Paschmann et al., 2013). Discontinuity detection is not very difficult, especially using algorithms based on angular changes. In contrast, distinguishing between classical TDs and RDs is difficult. Recent studies have approached discontinuity classification from a rather different perspective. Greco et al. (2016), for example, classified discontinuities in terms of their internal structure: (a) those with simple transitions from one side to the other are referred to as isolated events, and (b) those associated with complex networks of multiple small-scale interconnected discontinuities, are referred to as connected events.

Interplanetary discontinuities arrive at 1 AU and interact with the Earth's bow shock. Karlson et al. (2022) showed how some DDs can pass through the bowshock almost unchanged, while Kropotina et al. (2021) argued that the interaction with the Earth's bow shock can significantly alter discontinuity structure and stability. Webster et al. (2021) studied the interaction between solar wind discontinuities and the Earth's bow shock, and showed that discontinuities become thinner and that their current density (a measure of their strength) increases in the magnetosheath.

The complex nature of the solar-terrestrial system imposes more advanced tools to be used in computational space physics. In recent years, there has been a clear growth of published articles on applied machine learning techniques in space plasmas, such as solar wind characterization and prediction (Li et al., 2020; Upendran et al., 2020), space whether research (Camporeale et al., 2018; Camporeale, 2019), forecasting radiation belt dynamics (Bernoux et al., 2021), and geomagnetic storm prediction (Cristoforetti et al., 2022). Machine learning (ML) algorithms can be used to build models based on a training data set, and then try to make predictions without being explicitly programmed how to do so. In this study we use a hybrid method, based on convolutional neural networks (CNN) and support vector machines (SVM), for a binary classification of interplanetary discontinuities.

Munteanu et al. (2022) describe a hardware, field programmable gate-array (FPGA), implementation of a discontinuity detector, designed for use on-board a satellite to continuously monitor local magnetic field rotation angles. A software implementation of that discontinuity detector is included in the freely-distributed software analysis tool called Integrated Nonlinear Analysis (INA; Munteanu et al., 2023; see also the PhD thesis Munteanu, 2017). In this study we further develop this discontinuity detector by designing and implementing a novel multiscale detection and classification algorithm for discontinuities. This improved algorithm can automatically detect and classify discontinuities, based on classification criteria similar to those in Greco et al. (2016). Localized time-scale images depicting angular changes for each event are created, and then used as input for supervised machine learning classification schemes. In-situ magnetic field observations in 2007 and 2008 from ESA's Cluster mission are used to test and validate our detection and classification approach.

The paper is structured as follows. Section 2 presents the in-situ magnetic field observations used in our study. Section 3 describes the discontinuity identification algorithm, and presents the catalogue of events. Section 4 introduces the preliminary classification scheme, and presents the supervised machine learning models. Section 5 shows the results and discusses the accuracy of the CNN-SVM classifier. Our investigation of the occurrence rate of interplanetary discontinuities as a function of spacecraft location, is also included in this section. We give our conclusions in Section 6.

## 2. Data

We use in-situ observations from ESA's Cluster mission in 2007 and 2008, a multi-spacecraft mission with nearly 90° inclination elliptical polar orbit, perigee at about 4 $R_E$ and apogee at about 20 $R_E$ geocentric distance (1 $R_E$ = 6371 km), and an orbital period of approximately 57 h (Escoubet et al, 2001). Cluster enters the upstream solar wind during apogee in January-April every year, therefore we focus only on these intervals. We use spin resolution (4 s) magnetic field measurements from the fluxgate magnetometer on-board Cluster 1 (C1) spacecraft (Balogh et al., 2001; https://cdaweb.gsfc.nasa.gov/cgi-bin/eval2.cgi?dataset=C1_CP_FGM_SPIN&index=sp_phys).

Individual orbits centered on perigee are extracted. An example is depicted in Figure 1, which shows orbit no. 5 of the 2008 data set. The C1 spacecraft is in the solar wind around apogee at 9:05 UT on January 12. Typical solar wind observations are characterized by relatively small-amplitude magnetic field fluctuations and an average field magnitude below ~10 nT. As the spacecraft approaches Earth, it will cross the bow shock and enter into the magnetosheath, which is characterized by larger-amplitude fluctuations. As the spacecraft moves even closer to Earth, it will cross the magnetopause and enter the magnetosphere. At orbit perigee the magnetic field magnitude attains a maximum value of ~1500 nT, and then decreases as the spacecraft moves away from the Earth, again encountering the magnetosheath and the solar wind.

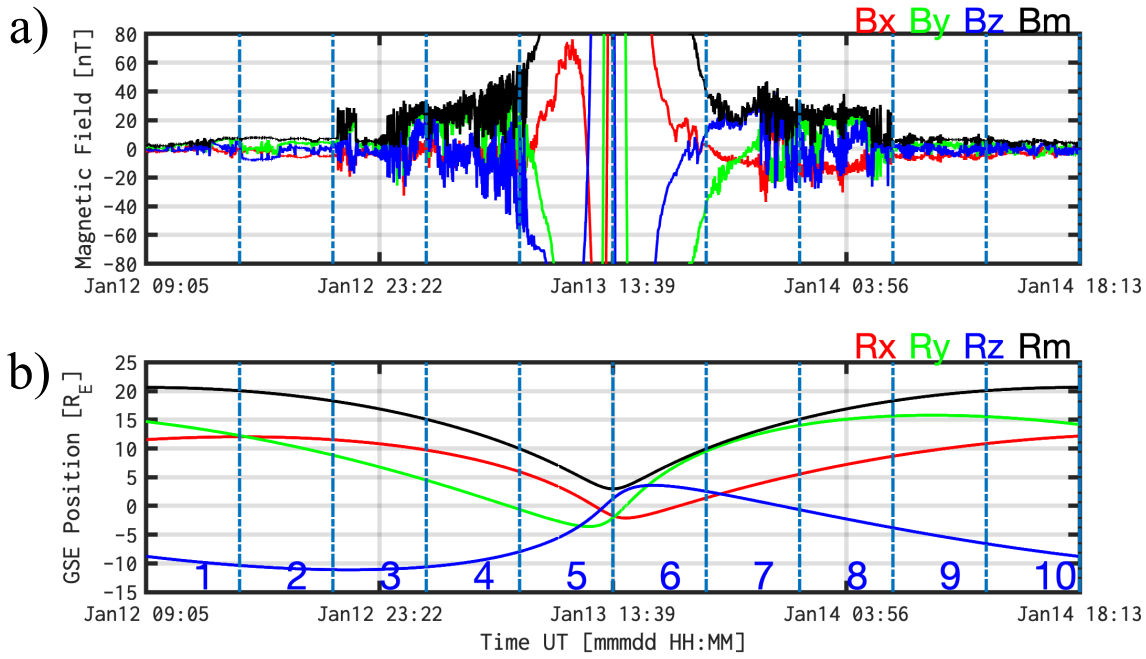


Figure 1. Illustration of data selection methodology. Individual orbits centered on perigee, are exctracted; shown is orbit no. 5 in 2008. Top panel: GSE magnetic field from Cluster 1 (C1) spacecraft; Bx, By, Bz, and magnitude Bm, are depicted using red, green, blue and black, respectively. Bottom panel: GSE position of C1. Numbers from 1 to 10 mark sub-intervals (see text for details).

Each individual orbit is further divided into 10 smaller intervals, of about 5.7 h each, labeled in Fig. 1 using numbers from 1 to 10. Based on Fig. 1, we assume that intervals labeled 1, 2, 9 and 10, correspond to unperturbed solar wind regions; intervals labeled 3, 4, 7 and 8, contain magnetosheath observations; and during intervals 5 and 6, the spacecraft is

inside the Earth's magnetosphere. The division into orbits and intervals is designed to: (a) manage the computation time and computer resources required to generate the results, and (b) obtain a way of estimating the discontinuity occurrence rate as a function of spacecraft location.

Following the example in Fig. 1, each data set is divided into individual orbits. The result is a total number of 100 orbits distributed equally among the two years, that is, a set 50 orbits for January-April 2007 and another set of 50 orbits for January-April 2008.

## 3. Identification

Let us consider the magnetic field vector $\boldsymbol{B}(\boldsymbol{t}) = [B_x\ B_y\ B_z]$, in an arbitrary reference system. Magnetic field directional discontinuities are characterized by sharp changes in the direction of this vector, computed as:

$$\varphi(t_k) = \left(\frac{180}{\pi}\right) \cos^{-1}\left(\frac{\boldsymbol{B_1}\cdot\boldsymbol{B_2}}{|\boldsymbol{B_1}|\cdot|\boldsymbol{B_2}|}\right) \quad (1)$$

where $\varphi$, in degrees, is computed at time $t_k$; $\boldsymbol{B_1} = \langle[B_x\ B_y\ B_z]\rangle_{\tau_1}$ and $\boldsymbol{B_2} = \langle[B_x\ B_y\ B_z]\rangle_{\tau_2}$, with the symbol $\langle\cdot\rangle_\tau$ denoting time averaging. We define a window $W$ centered on time $t_k$: $W = [t_{k-\tau/2}, t_{k+\tau/2}]$, with $\tau$ denoting the length of this window. Relative to $t_k$, the intervals $\tau_1$ and $\tau_2$ are defined as: $\tau_1 = [t_{k-\tau/2},\ t_k]$ and $\tau_2 = [t_k,\ t_{k+\tau/2}]$. Clearly, $\tau_1$ and $\tau_2$ contain the same number of data samples ($\tau_1 = \tau_2 = \tau/2$); in the following, we will refer to $\tau$ as the analysis scale. According to Equation (1), $\varphi$ takes values between 0° (parallel orientation between $\boldsymbol{B_1}$ and $\boldsymbol{B_2}$) and 180° (antiparallel orientation).

For the analysis of a signal continuously sampled in-situ, we developed a sliding-window algorithm which computes the angular changes $\varphi$ for windows $W$ centered at each time instance $t_k$. In the case of an isolated discontinuity, a specific variation of angular changes is artificially created by this sliding window algorithm. Due to the relative position of the sliding window with respect to the center of the discontinuity, an increasing (decreasing) trend in angular changes results, as the window moves closer to (away from) the discontinuity center; $\varphi$ attains a maximum value at the center of an isolated discontinuity.

The discontinuity detection algorithm is based on a critical value of the angular change. This value, denoted as $\varphi_c$, is set here to 30°. We define a local discontinuity measure (*LDM*), and use it as a quantitative measure for the presence of directional discontinuities. *LDM* is defined to be equal to $\varphi$, when $\varphi \geq \varphi_c$, and zero otherwise:

$$LDM(t_k) = \begin{cases} \varphi(t_k),\ if\ \varphi(t_k) \geq \varphi_c \\ 0,\ otherwise \end{cases} \quad (2)$$

Described above is the discontinuity detector as it was implemented by Munteanu et al. (2022), and included in the software tool INA (Munteanu et al., 2023). In this study we further develop this algorithm, by computing a matrix $LDM(t, \tau)$, defined for a series of scales $\tau$. As a quantitative measure for the presence of discontinuities, in this updated version of the discontinuity detector we inspect the values of $LDM^m(t)$, defined as the average value of $LDM(t_k)$ over a range of scales $\tau$. The averaging procedure is designed to minimize the effect random fluctuations, and enable the identification of only the dominant discontinuities.

Figure 2 illustrates our new identification procedure, applied to interval labeled 1 in Fig. 1. Figure 2a shows the magnetic field observations during this interval. Multiple abrupt changes in the filed component amplitudes are observed, as, for example, the jump in the Bz

component from -2 nT to +2 nT, observed around sample number 1100. Figure 2b depicts the measure $LDM^m(t_k)$ for this interval, and it is clear that most of the abrupt amplitude changes observed in Fig 2a correspond to values of $LDM^m > 30°$, thus, they are catalogued as events.

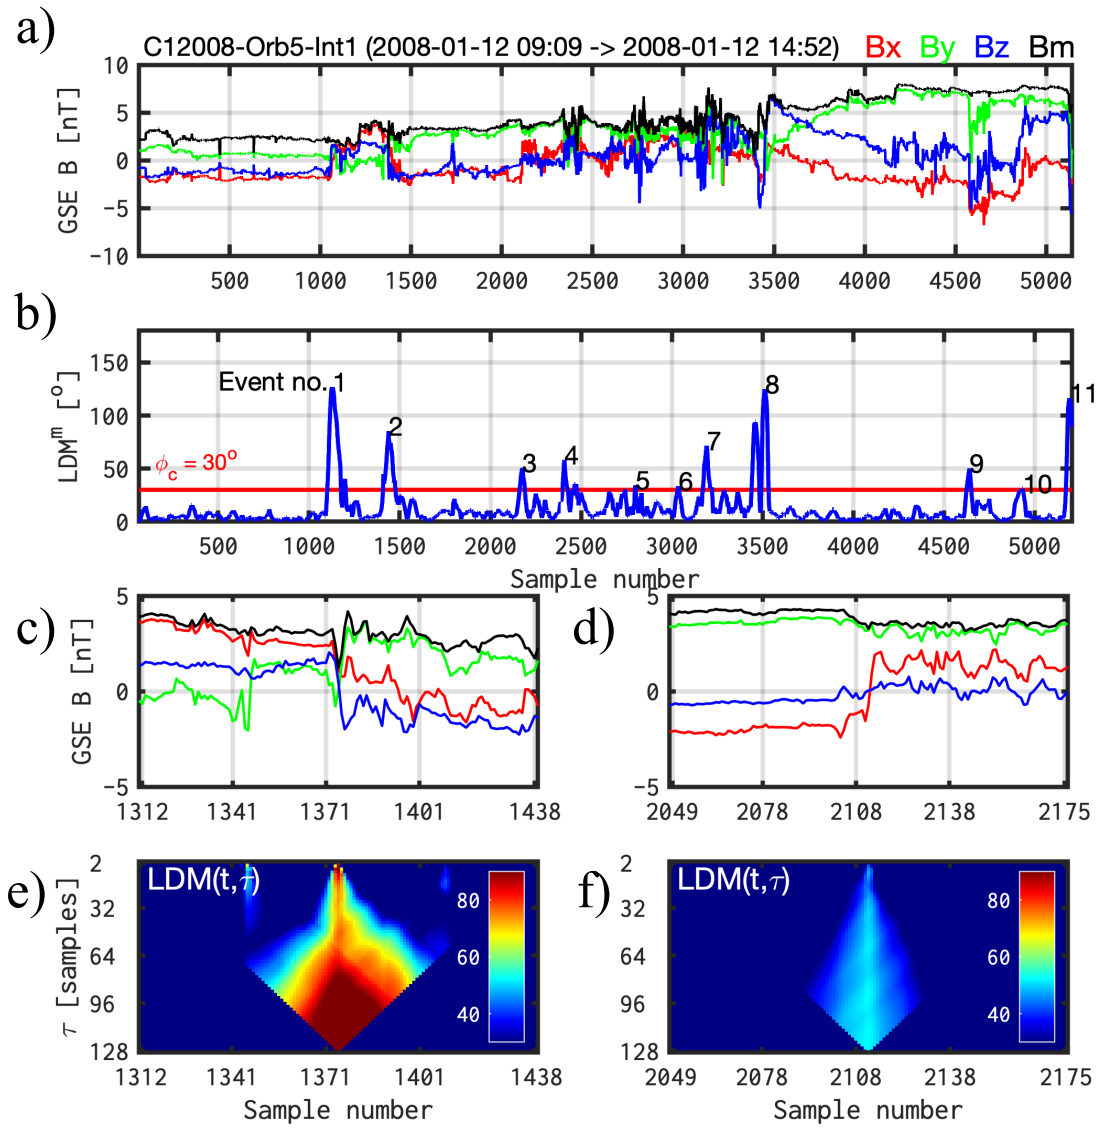


Figure 2. Illustration of directional discontinuity (DD) identification procedure. Shown is interval no. 1 of orbit no. 5 in 2008. Panel a): magnetic field versus sample number. Panel b): blue line shows $LDM^m$ versus sample number; red line marks the threshold $\varphi_c = 30°$. Panels c) and e): magnetic field for 128 samples centered on event no. 2, and the corresponding color representation of the matrix $LDM(t, \tau)$. Panels d) and f): similar to c) and e), but for event no. 3.

A number of 11 directional discontinuities are identified within this interval, with corresponding angular changes ranging from 30° (weak events) to 130° (strong events), see Fig. 2b. We note that the discontinuities identified here are not equally distributed in time, that is, periods with almost no discontinuities are followed by periods with multiple discontinuities. This is a well-known results; Burlaga (1969) was among the first to show that

there is a tendency for discontinuities to cluster together. Since then, many other studies have revealed this property (see, e. g., the review by Tsurutani et al., 1999).

Figures 2c and 2d depict the magnetic field observations in a window of 128 samples centered on events no. 2 and 3, respectively. Corresponding time-scale, color representations of $LDM(t,\tau)$ for the two events are depicted in Figs. 2e and 2f, respectively. Figure 2b showed that the maximum value of $LDM^m$ is around 90° for event no. 2, and around 50° for event no. 3; from this, we can state that event no. 2 is stronger than event no. 3. In Fig. 2e, two weak discontinuities can be observed on each side of the central event: one around sample no. 1341 and the second one around sample no. 1401; their strength decreases rapidly as we go to larger scales. Since our algorithm is designed to select only the strongest discontinuity within the analyzed window, only the central DD is automatically identified and catalogued by our code. In contrast, DD no. 3 is clearly isolated from other structures (see Fig. 2f). This discussion is relevant for discontinuity classification, as shown in the next section.

The identification procedure illustrated in Figure 2 is applied to our data set of 100 orbits in 2007 and 2008. The final catalogue consists in 4215 events in 2007 and 5194 in 2008.

## 4. Classification

### 4.1. Preliminary Classification

The previous section showed two examples of events: (a) one with a rather complex local environment, where the central discontinuity was flanked by two weaker events (Fig. 2e); and (b) a clean discontinuity, clearly separated from other structures or fluctuations (Fig. 2f). Greco et al. (2016) used a similar reasoning to classify their set of event as "connected" and "isolated". Inspired by this, and also by results from Fig. 2, these two discontinuities are chosen as representatives for our pre-classification scheme: events resembling that in Fig 2e will be called "complex", and events resembling that in Fig. 2f, will be called "simple".

With almost 10000 events in our database, a visual-based classification is clearly not possible. An automated classification algorithm was designed to distinguish between the two classes. The current version of the discontinuity detector uses a number of $ns = 64$ scales. For each scale, local peaks in $LDM(t)$ (defined in Equation 2) are found using the MATLAB built-in function called "findpeaks" (https://www.mathworks.com/help/signal/ref/findpeaks.html). The result is $np(\tau)$, the number of peaks per scale. A "complexity index" CI is defined as:

$$CI = \frac{1}{ns}\sum_{\tau} np(\tau), \quad (3)$$

As a rule of thumb, if CI has a high value, then the discontinuity is classified as complex; if it has a small value, then it is classified as simple. By trial and error, a threshold value of *CI=1.1* was found to provide good results. Although this choice seems arbitrary, it is confirmed by our extensive validation tests presented in Section 5.

The preliminary classification procedure is illustrated in Figure 3, which shows the number of peaks per scale, $np(\tau)$, and the complexity index *CI*, for the set of 11 events identified in Fig. 2b. Fig 3a shows that events nos. 3, 6, 10 and 11, have one peak per scale, meaning that they are ideal simple events. Events nos. 1, 2, 4, 5, and 7 are complex, because

multiple peaks per scale are found. The complexity index depicted in Fig. 3b confirms our interpretation based on the values of $np(\tau)$.

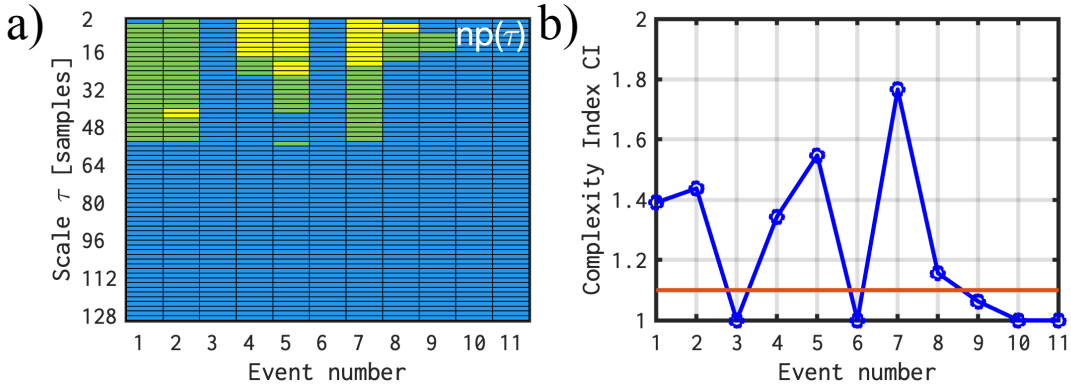


Figure 3. Illustration of the preliminary classification procedure. Shown are the 11 events identified in interval no. 1 of orbit no. 5 in 2008. Panel a): number of peaks $np(\tau)$, color coded, versus scale (on the y-axis) and event number (on the x-axis); blue corresponds to $np = 1$, green depicts $np = 2$, and yellow signifies that $np \geq 3$. Panel b): complexity index CI versus event number; red line marks the boundary between simple and complex events, at $CI = 1$.

Figure 3b also shows that two events are very close to the (somewhat arbitrarily chosen) threshold between simple and complex: events nos. 8 and 9. Technically, event no. 8 is pre-classified as complex and event no. 9 is pre-classified as simple, but, because they are so close to the threshold, they are referred to as "mixed" events. The next section discusses how this type of mixed events represent the main source for the differences between pattern recognition and pre-classification results.

Figure 2 showed LDM images with the full range of angular changes, and we noted that discontinuities can vary in strength from 30° (weak events) to 180° (strong events). This introduces additional variability in the observed patterns. This adverse effect can be minimized by setting to 0 all angular changes below 30°, and to 1, all those above 30°. Thus, a new set of LDM "binary" images is generated, where all angular changes below 30° are depicted in black, and those above 30° are shown with white color. Examples from this re-processed set of LDM images, which will be used to train the machine learning models, are shown in Figure 4, which depicts the 11 events identified in Fig. 2b, and pre-classified in Fig. 3b. For most cases, a clear distinction is observed between events pre-classified as simple and those pre-classified as complex. Events nos. 3, 6, and 9, are correctly pre-classified as simple, because no other structures are observed outside of the central region. The diamond shape of event no. 3, for example, is what we ideally expect for a simple discontinuity.

All pre-classified complex events in Fig. 4 have visible structures near or overlapping with the central region. Event no. 8 has a complexity index very close to the threshold value, but it is technically pre-classified as complex. In the next section we will show that the machine learning algorithm based on pattern recognition, predicts this event as simple.

The preliminary classification procedure described above was applied to the 10 intervals of orbit no. 5. The results are shown in Figure 5. As Cluster 1 spacecraft orbits around Earth, it will cross the solar wind (intervals 1 and 2), the magnetosheath (intervals 3 and 4) and the magnetosphere (intervals 5 and 6) then back through the magnetosheath (intervals 7 and 8) and solar wind (intervals 9 and 10).

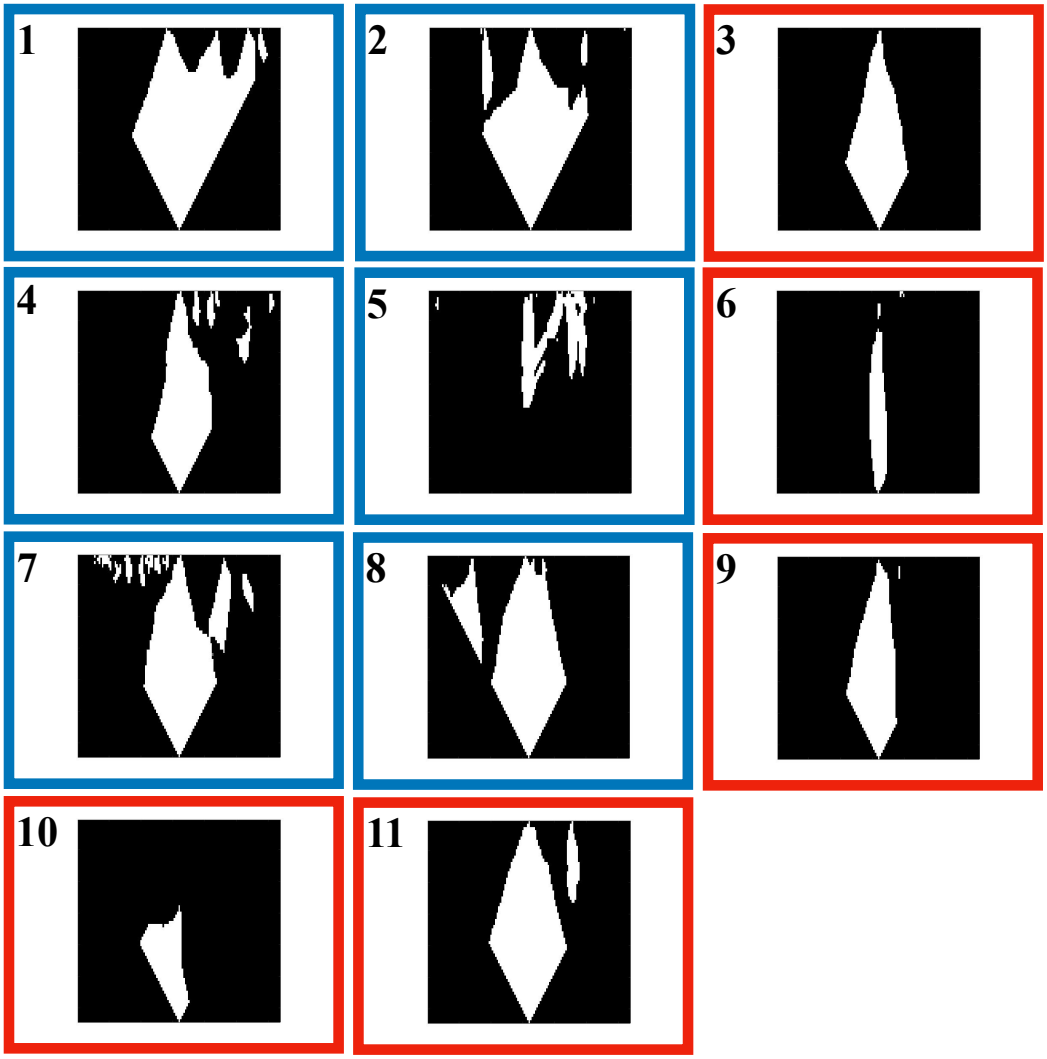


Figure 4. LDM binary images for the 11 events from Fig. 3. Event number is in the top-left corner of each image, and color indicates the event class: blue is complex and red is simple.

Fig. 5a shows the event occurrence rates, expressed in number of events per hour, and Fig. 5b shows the corresponding ratio between the numbers of simple and complex events, expressed in percentages. The illustration in Fig. 5a allows one to investigate the variation of event rate as a function of spacecraft location. Fig. 5a shows that the average occurrence rate is around 2 DDs/h in both the solar wind and the Earth's magnetosheath region, and that virtually a number of 0 DDs/h are identified inside the magnetosphere. Although the total rate remains constant as we cross from the solar wind to the magnetosheath interval, the rates and the corresponding percentages of simple/complex events change considerably: intervals 3 and 4 have a rate of occurrence of ~0.5 DDs/h for simple, and ~1.5 DDs/h for complex events. The corresponding percentages are 30/70 %, for simple/complex, respectively.

The results for the second half of the orbit depicted in Fig. 5b look somewhat different compared to the first half. A higher than expected percentage of solar wind events are pre-classified as complex. Three possible sources of variability can affect our results: (a) large data gaps, (b) solar wind variability and (c) possible asymmetry between the left-sides and the right-sides of each orbit. A dedicated study of these effects is necessary, but it is outside the scope of this study.

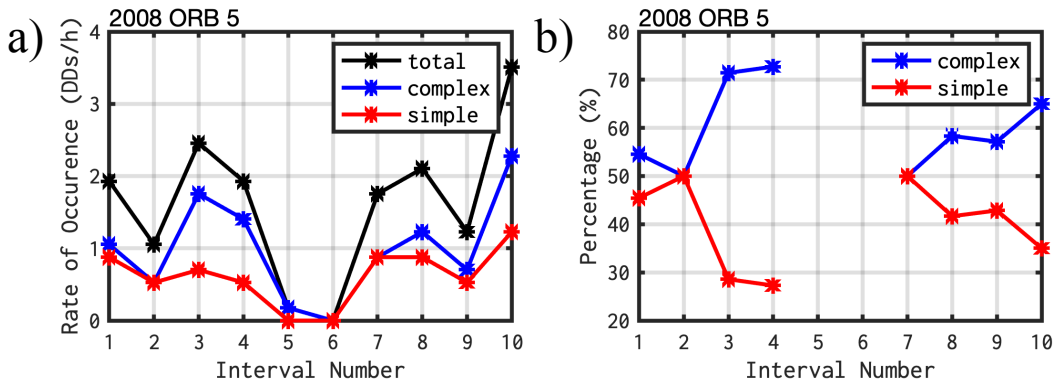


Figure 5. Rate of occurence versus interval number, for orbit no. 5 in 2008. Panel a): rate of occurrence expressed in number of events per hour; total rate is in black, complex is in blue and the occurrence rate of simple events is in red. Panel b): percentages of simple and complex events, relative to the total number of events.

All ~10000 events in our catalogue were pre-classified as described above. In January-April 2007, a number of 1806 events are pre-classified as simple and 2410 as complex. In January-April 2008, 1997 events are simple and 3197 are complex. LDM binary images are created for each event, and will be used in the next section to train neural networks for pattern recognition.

## 4.2 Classification using supervised machine learning

There are two main categories of machine learning models: supervised and unsupervised. Supervised learning uses labelled data, known as training data, to learn a specific pattern. On the other hand, unsupervised learning uses unlabeled data, and can reveal unanticipated patterns and relationships. Here, we adopt a supervised learning approach with the objective of creating models that provide accurate predictions of the pre-classified LDM images described in the previous section.

We propose a hybrid convolutional neural network (CNN) and support vector machines (SVM) which uses the feature extraction capability of CNNs, and combines it with the powerful classification features of SVMs (Cortes and Vapnick, 1995) for a binary classification problem. We decided to train two CNN-SVM models. One model is trained on the 2007 dataset and is used to predict the events in 2008 and, vice-versa, the predictions for 2007 are done using a model trained on 2008 events. Figure 6 illustrates our workflow. Some technical details are provided below; more details on network architectures can be found in Bishop (2006), or LeCun et al. (2019).

The pre-classified LDM images constitute the labelled data used to train the neural networks. From all pre-classified images for each class, 70 % is the training set and 30% is the test set, separately for each year. The training set is considered ground-truth, and is used to update all the parameters in the training step. All images are first contracted to 64 x 64 pixels in order to speed-up the processing time.

CNNs are traditionally used for pattern recognition due to their ability to extract features with a high degree of abstraction. Our models follow a typical CNN architecture consisting of two convolutional modules and a fully connected layer, that are stacked on top of each other. Each module includes a convolutional layer followed by a pooling layer; a

ReLU activation function was used after each stage. The convolutional layer acts as feature extractor, and in our models consists of 32 filters, extracting 3 × 3 pixel subregions. The Pooling layer is applied after the convolutional layer in order to reduce the spatial resolution; a max-pooling method was used, with $size = 2$ and $stride = 1$. The feature map obtained from the convolutional operations is flattened in a 1-D vector and feeds the last layer composed of 128 neurons fully connected. The cost function we choose to minimize is the hinge-loss, which is used for "maximum-margin" classification in SVM. In order to avoid overfitting, an L2 regularization or Ridge regression was implemented in the training process.

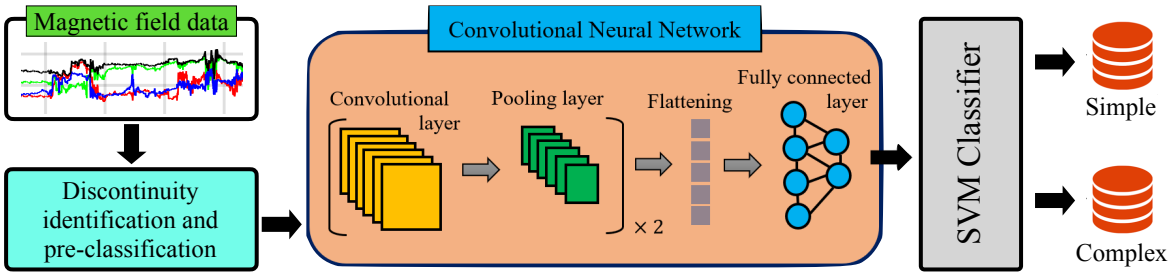


Figure 6. Supervised machine learning workflow. From magnetic field data we identify discontinuities and create the LDM binary images which are then used as input for CNN. The SVM classifier predicts the class of each event.

In the training process, the pre-classified LDM images are considered as ground-truth labels which will update all the network's parameters, and the loss function is calculated. The process is repeated for a given number of epochs until the loss function reaches a minimum, and, if it does not improve, the network training is halted. An Adam optimization algorithm was selected to update model parameters. The training and testing of our models was done in Python 3.9 with TensorFlow libraries on a standard PC, on a single NVIDIA Quadro 5000 graphics card.

## 5. Validation

### 5.1 Confusion matrices and derived model performance metrics

We use confusion matrices to evaluate the performance of our models and to visualize classification results. For a binary classification, the confusion matrix is a 2 × 2 contingency table. Figure 7 shows the confusion matrices for 2007 and 2008. In each case, the main diagonal shows the number of correctly predicted events: true positive (TP; top left), which, for our case, is this the number of pre-classified simple events that are predicted correctly; true negative (TN; bottom right), is the number of pre-classified complex events that are predicted correctly. The off-diagonals of each matrix shows the number of incorrectly predicted events: false negative (FN; top right), also known are type I error, is the number of pre-classified simple events that are incorrectly predicted as complex; false positive (FP; bottom left), also known as type II errors, is the number of pre-classified complex events that are incorrectly predicted as simple. The sum TP+FN is the total number of pre-classified simple events, and TN+FP is the total number of pre-classified complex events.

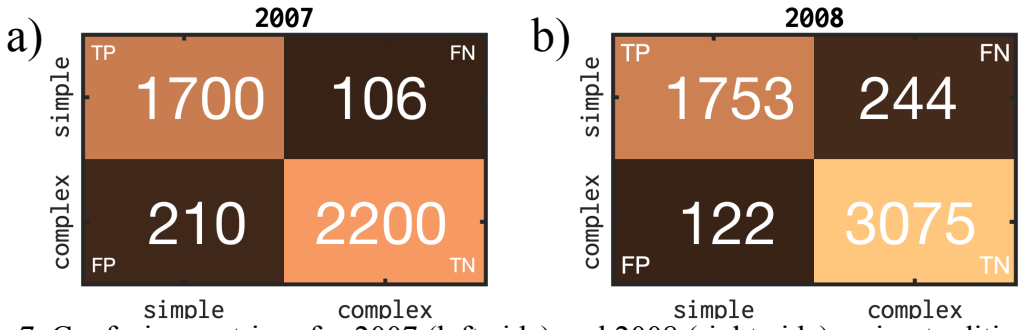


Figure 7. Confusion matrices for 2007 (left-side) and 2008 (right-side), using traditional notations in a binary contingency table. The number of events predicted for each class are shown in each cell.

Figure 7a shows the confusion matrix for 2007. From a number of 1806 pre-classified simple events, the model predicts only a number of TP=1700 events as simple and the rest are predicted as complex (FN=106). On the complex side, from 2410 pre-classified complex events, the model predicts only a number of TN=2200 events as complex, and the rest are predicted as simple (FP=210). Figure 7b shows the corresponding results for 2008. From a number of 1997 pre-classified simple events, the model predicts only a number of TP=1753 events as simple and the rest are predicted as complex (FN=244). From 3197 pre-classified complex events, the model predicts that only a number of TN=3075 events are complex, and the rest are predicted to be simple (FP=122).

Table 1. Common metrics used to evaluate a model's performance.

| Metric name | Definition | Results for 2007 dataset | Results for 2008 dataset |
|---|---|---|---|
| Accuracy (Acc) | $\frac{TP+TN}{TP+TN+FP+FN}$ | 0.925 | 0.929 |
| Precision (Pre) | $\frac{TP}{TP+FP}$ | 0.890 | 0.935 |
| Recall (Rec) | $\frac{TP}{TP+FN}$ | 0.941 | 0.877 |
| F1 *score* | $2\times\frac{Pre\times Rec}{Pre+Rec}$ | 0.914 | 0.905 |
| MCC | $\frac{(TP\cdot TN)-(FP\cdot FN)}{\sqrt{(TP+FP)(TP+FN)(TN+FP)(TN+FN)}}$ | 0.849 | 0.850 |

The information contained in a confusion matrix can be used to derive some of the most common metrics used to evaluate a model's performance. Table 1 gives the definitions for some of these metrics, using the standard notations from a 2 × 2 confusion matrix. Acuraccy is the ratio between the number of correct predictions and the total number of predictions. Precision is the ratio of the correctly predicted simple and the total number of predicted simple. Recall is the ratio of correctly predicted simple divided by the number of pre-classified simple events. If we optimize for Recall, it will decrease FN (incorrectly predicted complex events) and increase TP with the cost of increasing FP (the number of incorrectly predicted simple events). Due to their nature, Precision and Recall are always in a mutual trade-off relationship. The *F1 score* quantifies the model's ability to predict both

classes correctly, based on the harmonic mean of Precision and Recall. Hence, if either Precision or Recall has a low value, the F1 score suffers significantly (Powers, 2020).

One of the most popular choices for estimating a model's performance is the Matthews Correlation Coefficient (MCC). This measure is more informative than the F1 score because it takes into account the balance ratios of all four confusion matrix cells (Chicco and Jurman, 2020). MCC can have a minimum value equal to -1, indicating a complete disagreement between pre-classification and prediction, and a maximum value of 1, indicating a perfect prediction accuracy.

Table 1 also shows the values for the different metrics. Almost all values are close to 0.9, meaning that the model predictions are accurate. The values for MCC are slightly below 0.9, but this does not necessarily mean that the predictions are not accurate. Our preliminary criteria to distinguish between simple and complex are somewhat arbitrary. In Section 4.1 we mentioned that some events are "mixed", and we argue that these are the main cause for the differences between preliminary classification and machine learning prediction. Thus, even though the values for MCC might suggest that the prediction is not perfect, some of the events predicted "incorrectly" by the ML models might in fact be better classified than the preliminary classification. This statement is supported by the machine learning predictions for the images depicted in Fig. 4. ML prediction is in agreement with the pre-classification for almost all events, except for event nos. 8 and 11. Event no. 11, for example, was pre-classified as simple, but the ML model classified it as complex. Comparing the LDM pattern observed for event no. 11 with that for event no. 3, it is rather obvious that event 11 is better classified as complex.

We showed above that the differences between machine learning predictions and pre-classification results are rather small. Thus, in the following, we use only the machine learning predictions to investigate the rate of occurrence of interplanetary discontinuities as a function of spacecraft location.

## 5.2. Rate of occurrence of interplanetary discontinuities

Figure 5 showed that the total rate of occurrence for the events in orbit no. 5 remained approximately constant, at about 2 DDs/h, as the spacecraft crossed from the solar wind into the magnetosheath; inside the Earth's magnetosphere the total occurrence rate was close to zero. Figure 5 also showed that the percentages of simple/complex events was close to 50/50 % in the solar wind, but, in the magnetosheath, the number of complex events increased significantly, reaching a value of 70 % from the total number of events. Figure 8 shows the results for the entire set of orbits. Same as before, 2007 and 2008 are investigated separately, and then compared to each other.

Let us first consider the variability of the total rate of occurrence for the two data sets (Figs. 8a and 8b). The upper envelope varies between 4 and 5 DDs/h, for both years. The lower envelope is 0 DDs/h; most likely, this is the result of data gaps, but, since their number is relatively small, their effect is correspondingly small. The variability of the total rate of occurrence is related to the variability of the solar wind itself. For example, it is well known that fast solar wind contains more discontinuities than the slow solar wind (see, e.g., Section 4.1 in the paper by Borovsky et al., 2010).

Let us now consider the variability of the rate of occurrence for each class separately. The upper envelope for the complex events in 2007 (Fig. 8c), starts at 2.5 DDs/h in the solar wind, increases to about 5 DDs/h, and then decreases towards 0 for intervals 5 and 6. A similar variation can be seen for the second part of the orbit. In 2008 (Fig. 8d), the upper

envelope for the rate of occurrence of complex events follows closely that for the total rate, implying that most of the variability of the total rate in 2008 comes from complex events. The upper envelope for simple events (Figs. 8e and 8f) is around 2 DDs/h for all solar wind and magnetosheath intervals, in both years, confirming that the variability of the total rate of occurrence is dominated by that of complex events.

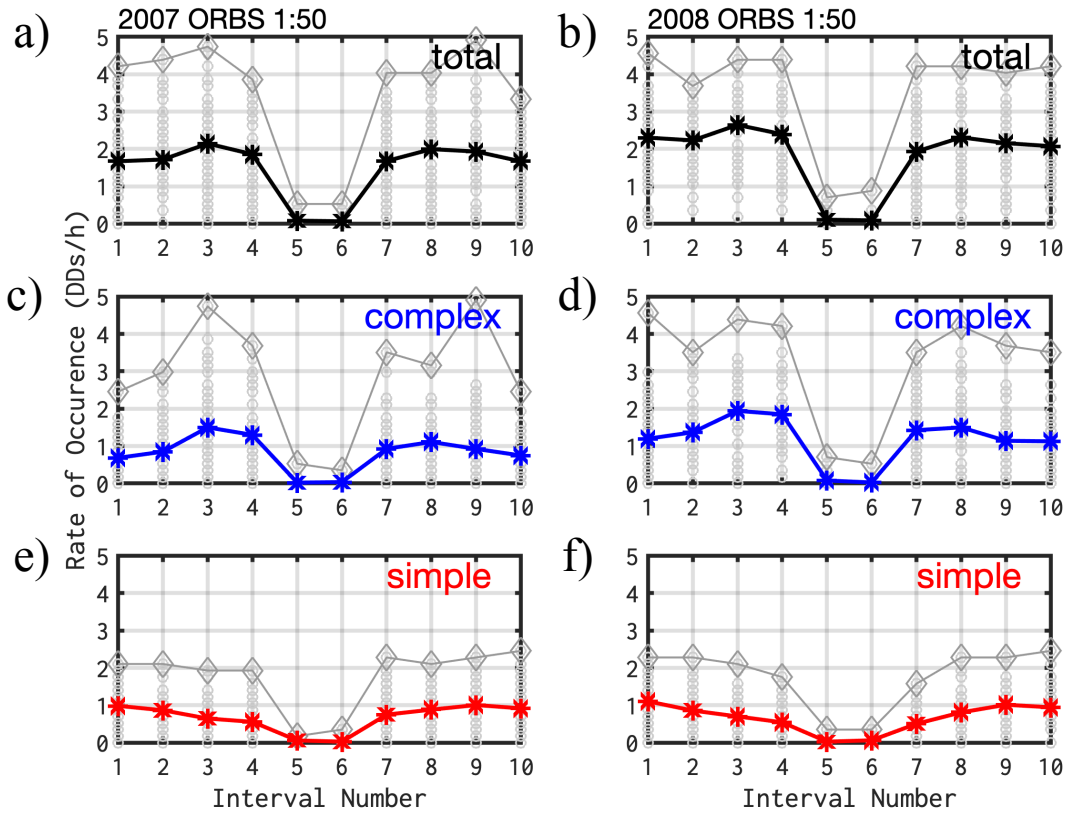


Figure 8. Rate of occurrence expressed as number of events per hour, versus interval number. Left-column: results for the set of 50 orbits of C1 spacecraft in January-April 2007. Right-column: corresponding results in 2008. Top: total rate of occurrence; middle: the rate for complex events; bottom: the rate of occurrence for simple events. In each panel: grey circles mark each orbit; grey diamonds mark the upper envelope for each interval; and thick lines depict mean values in each set.

Figure 8 also shows the mean values for each set. The mean value for the total rate is around 2 DDs/h for all solar wind and magnetosheath intervals, in both years. This implies that the total number of events does not change significantly as the spacecraft crosses from the solar wind into the magnetosheath. In other words, most solar discontinuities pass through the Earth's bow shock. The mean rate of occurrence for the complex events in 2007 increases from ~1 DDs/h in the solar wind to 1.5 DDs/h in the magnetosheath, it is 0 DDs/h inside the magnetosphere, and then the pattern reverses for intervals from 7 to 10 (Fig. 8c). A similar result is also observed in 2008 (Fig. 8d). Since the total mean rate remains approximately constant, but the mean rate for complex events increases inside the magnetosheath, we expect the mean rate for simple events to decrease inside the magnetosheath. This is exactly what we observe in Figs. 8e and 8f. The mean rate of simple events is around 1 DDs/h in the solar wind, decreases to around 0.5 DDs/h in the magnetosheath, is 0 inside the magnetosphere, and, as expected, the pattern reverses for intervals from 7 to 10.

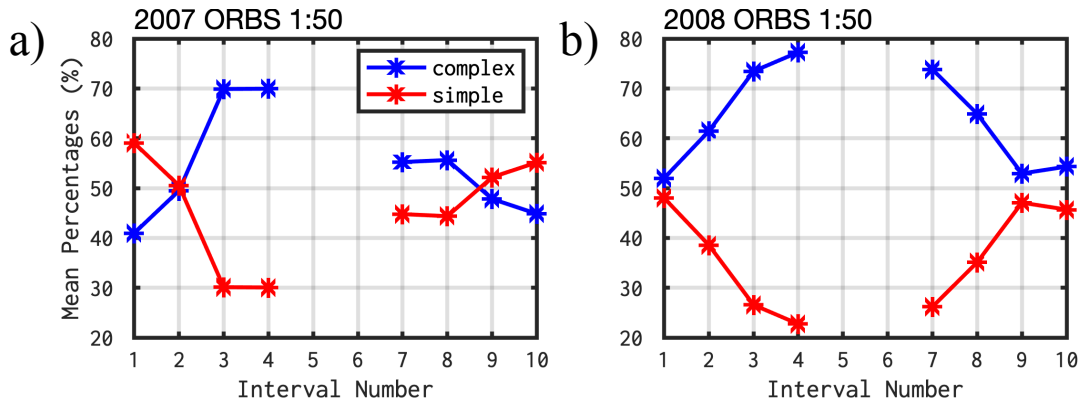


Figure 9. Mean percentages of predicted events for each class, relative to the total number of events, for 2007 (left-side) and 2008 (right-side). As before, complex events are shown in blue, and simple in red.

Figure 9 depicts the mean percentages of simple/complex events. In 2007 (Fig. 9a), interval 1 has 40/60 % mean percentages for simple/complex events, and the mean percentages are exactly 50/50 % for interval 2. The percentage of complex events increases in the magnetosheath to 70 % of the total, while the percentage of simple events correspondingly decreases to 30 %. In 2008 (Fig. 9b), the percentage of complex events increases monotonically, from ~50 % in interval 1, to almost 80 % in interval 4. For the second part of the orbit, a similar pattern is observed, but somewhat distorted. This is most probably an orbital effect. Further investigation of this orbital asymmetry is outside the scope of our study.

To our knowledge, Greco et al. (2016) is the only study that classified solar wind discontinuities using criteria similar to ours: they describe “connected” events, resembling our complex ones, and “isolated” events, resembling our simple ones. From the analysis of a 2 h interval of high resolution data from Cluster 4 spacecraft on 20 January 2007, they identified 1245 small-scale solar wind discontinuities, and determined a percentage of about 50 % between connected and isolated events. As discussed above, this is an almost identical percentage to that determined by us. This is a clear confirmation of our approach, at least for the solar wind intervals, and deserves further exploration.

The results obtained by us for the magnetosheath intervals can be compared to those from Webster et al. (2021). They showed that discontinuties change as they cross from the solar wind into the magnetosheath, becoming narrower and at the same time surrounded by larger amplitude fluctuation. This is consistent with our results, because narrower discontinuties surrounded by large amplitude fluctuations resemble complex events. The result from Webster et al. (2021) implies that an initially simple event in the solar wind will be classified as complex after its passage in the magnetosheath.

## 6. Summary and Conclusions

We designed and implemented a novel identification algorithm for interplanetary directional discontinuities. We used magnetic field observations from the Cluster 1 (C1) spacecraft in orbit around Earth, to test and validate our results. The detection algorithm is based on identifying abrupt changes of the direction of the magnetic field, referred to as directional discontinuities (DDs). Using a sliding window approach, angular changes for each

data point are computed; we repeat with successively increasing window lengths, and the result is a matrix of angular changes. Next, mean angular changes are computed, by averaging over the set of scales. Finally, discontinuities are identified as localized peaks of the series of mean angular changes. This algorithm was applied to magnetic field data from C1 in January-April 2007 and January-April 2008. A number of 4216 events were indentified in 2007, and 5194 in 2008.

Our main goal was the development of supervised machine learning models able to classify the events. For this, we first had to design a pre-classification algorithm capable of creating the labeled data necessary to train the machine learning models. Our pre-classification approach is based on counting the number of local maxima of the matrix of angular changes, and then inspecting the average number of peaks per scale. By trial and error, we determined a specific threshold value for the average number of peaks, and: all events below threshold were pre-classified as simple, and those above threshold as complex. In 2007, 1806 events were pre-classified as simple and 2410 as complex. In 2008, 1997 events are simple and 3197 are complex.

Supervised machine learning is based on a pattern recognition approach, thus, it needs images as input. We generated time-scale images for each event, depicting with color the matrix of angular changes. A further step was necessary at this point: the full-color representation of LDM matrices introduces details that can distort the pattern recognition algorithm. Thus, for the final set of images, we used "binary" representations with simple black and white patterns, with black regions denoting angular changes below 30°, and white regions denoting angular changes above 30°. A machine learning tool was implemented from convolutional neural networks with the help of a support vector machines classifier. The pre-classified images were used to train the machine learning models. We created two ML models: one using the images from 2007, and a second one for 2008. The model trained using the images in 2007 was then used to classify the images in 2008; and vice-versa for 2008. We showed confusion matrices for the two years separately, and demonstrated that the differences between ML classification and pre-classification are rather small.

We validated our classification results by investigating the occurrence rate of events as a function of spacecraft locations. For this, we divided our data into individual obits centered on perigee. We extracted 100 orbits, distributed equally among the two years. Each orbit was further divided into a number of 10 equal intervals. This allowed us to investigate the dependence of our results on the plasma region traversed by the spacecraft: around apogee (interval nos. 1, 2, 9 and 10), we assume that the spacecraft is in the upstream solar wind; interval nos. 3, 4, 7 and 8 correspond to the magnetosheath; and during interval nos. 5 and 6 the spacecraft is inside the magnetosphere. By averaging results for each set of 50 orbits, we showed that the total rate of occurrence is rather constant, at about 2 DDs/h, for both solar wind and magnetosheath regions, in both 2007 and 2008.

We also showed that complex and simple events start with roughly equal occurrence rates in the solar wind, but, interestingly, the rate of complex events increases significantly in the magnetosheath. Since the total rate is constant, this means that part of the simple events in the solar wind are transforming into complex events into the magnetosheath. We quantified the difference by investigating their relative percentages. We showed that the percentage of complex events increases monotonically from a solar wind value of 50 % to almost 80 % in the magnetosheath. As expected, the number of simple events follows a reverse trend, that is, it decreases from 50 % in the solar wind to 20 % in the magnetosheath.

We demonstrate that our classification scheme can provide meaningful geophysical insights, and thus be relevant for future studies of interplanetary discontinuities. In future, we

plan to design more advanced classification schemes, using, for example, unsupervised machine learning algorithms.

## Acknowledgments

This work was supported by the Romanian Ministry of Research, Innovation and Digitalization under Romanian National Core Program LAPLAS VII – contract no. 30N/2023. The work of C. M. was supported by ESA PRODEX MISION, and National project PN-III-P1-1.1-TE-2021-0102.

## Open Research

We used spin resolution magnetic field measurements from the FGM instrument on-board Cluster 1 spacecraft (Balogh et al., 2001), available from: https://cdaweb.gsfc.nasa.gov/cgi-bin/eval2.cgi?dataset=C1_CP_FGM_SPIN&index=sp_phys. Some of our results were obtained using MATLAB; part of our computer codes were adapted from INA, a software application freely available from: http://www.storm-fp7.eu/index.php/data-analysis-tools. The training and testing of our machine learning models was done in Python 3.9 with scikit-learn (https://scikit-learn.org/stable/modules/classes.html) and TensorFlow (https://www.tensorflow.org) as main libraries, on a standard PC with a single NVIDIA Quadro 5000 graphics card. The catalogue of about 10000 LDM binary images generated by this research are available from: https://github.com/ISS-psm/ldm.

## References

Artemyev, A. V., Angelopoulos, V., Halekas, J. S., Vinogradov, A. A., Vasko, I. Y., & Zelenyi, L. M. (2018). Dynamics of intense currents in the solar wind. The Astrophysical Journal, 859(2), 95. https://doi.org/10.3847/1538-4357/aabe89

Artemyev, A. V., Angelopoulos, V., & Vasko, I. Y. (2019a). Kinetic properties of solar wind discontinuities at 1 AU observed by ARTEMIS. Journal of Geophysical Research: Space Physics, 124(6), 3858–3870. https://doi.org/10.1029/2019JA026597

Artemyev, A. V., Angelopoulos, V., Vasko, I. Y., Runov, A., Avanov, L. A., Giles, B. L., et al. (2019b). On the kinetic nature of solar wind discontinuities. Geophysical Research Letters, 46(3), 1185–1194. https://doi.org/10.1029/2018GL079906

Balogh, A., Carr, C. M., Acuña, M. H., Dunlop, M. W., Beek, T. J., Brown, P., et al. (2001). The cluster magnetic field investigation: Overview of in-flight performance and initial results. Annales Geophysicae, 19(10/12), 1207–1217. https://doi.org/10.5194/angeo-19-1207-2001

Bernoux, G. , A. Brunet, É. Buchlin, M. Janvier and A. Sicard (2021). An operational approach to forecast the Earth's radiation belts dynamics, , J. Space Weather Space Clim., 11 - 60 DOI: https://doi.org/10.1051/swsc/2021045

Bishop, C. M. (2006). Pattern Recognition and Machine Learning, Springer, New York. ISBN: 978-1-4939-3843-8

Borovsky, J. E. (2008), Flux tube texture of the solar wind: Strands of the magnetic carpet at 1 AU?, J. Geophys. Res., 113, A08110, doi:10.1029/2007JA012684.

620 Borovsky, J. E. (2010). Contribution of strong discontinuities to the power spectrum of the solar wind. Physical
621 Review Letters, 105(11), 111102. https://doi.org/10.1103/PhysRevLett.105.111102.

622 Borovsky, J. E., and M. H. Denton (2010), Solar wind turbulence and shear: A superposed-epoch analysis of
623 corotating interaction regions at 1 AU, J. Geophys. Res., 115, A10101, doi:10.1029/2009JA014966.

624 Burch, J. L., Moore, T. E., Torbert, R. B., & Giles, B. L. (2016). Magnetospheric multiscale overview and
625 science objectives. Space Science Reviews, 199(1–4), 5–21. https://doi.org/10.1007/s11214-015-0164-9.

626 Burkholder, B. L., & Otto, A. (2019). Magnetic reconnection of solar flux tubes and coronal reconnection
627 signatures in the solar wind at 1 AU. Journal of Geophysical Research: Space Physics, 124(11), 8227–8254.
628 https://doi.org/10.1029/2019JA027114.

629 Burlaga, L. F. (1969). Directional discontinuities in the interplanetary magnetic field. Solar Physics, 7(1), 54–
630 71. https://doi.org/10.1007/ BF00148406

631 Camporeale, E., S. Wing, J. Johnson (2018). Machine Learning Techniques for Space Weather, Elsevier
632 Science. ISBN: 9780128117897

633 Camporeale, E. (2019). The challenge of machine learning in space weather: Nowcasting and forecasting. Space
634 Weather, 17(8), 1166–1207. https://doi.org/10.1029/2018SW002061

635 Chian, A. C.-L., & Muñoz, P. R. (2011). Detection of current sheets and magnetic reconnections at the turbulent
636 leading edge of an interplanetary coronal mass ejection. The Astrophysical Journal Letters, 733(2), L34.
637 https://doi.org/10.1088/2041-8205/733/2/L34

638 Chicco, D., Jurman, G. (2020). The advantages of the Matthews correlation coefficient (MCC) over F1 score
639 and accuracy in binary classification evaluation. BMC Genomics **21**, 6. https://doi.org/10.1186/s12864-019-
640 6413-7

641 Cortes, C., Vapnik, V. (1995). Support-vector networks. Mach Learn **20**, 273–297.
642 https://doi.org/10.1007/BF00994018

643 Cristoforetti, M., Battiston, R., Gobbi, A. et al. (2022). Prominence of the training data preparation in
644 geomagnetic storm prediction using deep neural networks. Sci Rep 12, 7631. https://doi.org/10.1038/s41598-
645 022-11721-8

646 Escoubet, C. P., Fehringer, M., and Goldstein, M. (2001). Introduction: The Cluster mission, Ann. Geophys.,
647 19, 1197–1200, https://doi.org/10.5194/angeo-19-1197-2001 .

648 Greco, A., Chuychai, P., Matthaeus, W. H., Servidio, S., & Dmitruk, P. (2008). Intermittent MHD structures
649 and classical discontinuities. Geophysical Research Letters, 35(19), L19111.
650 https://doi.org/10.1029/2008GL035454

651 Greco, A., Matthaeus, W. H., Perri, S., Osman, K. T., Servidio, S., Wan, M., & Dmitruk, P. (2018). Partial
652 variance of increments method in solar wind observations and plasma simulations. Space Science Reviews, 214,
653 1. https://doi.org/10.1007/s11214-017-0435-8

654 Greco, A., Perri, S., Servidio, S., Yordanova, E., & Veltri, P. (2016). The complex structure of magnetic field
655 discontinuities in the turbulent solar wind. The Astrophysical Journal Letters, 823(2), L39.
656 https://doi.org/10.3847/2041-8205/823/2/L39

657 Greco, A., & Perry, S. (2014). Identification of high shears and compressive discontinuities in the inner
658 heliosphere. The Astrophysical Journal, 784(2), 163. https://doi.org/10.1088/0004-637X/784/2/163

659 Haaland, S., Munteanu, C., Mailyan, B., Pulkkinen, & Raststätter, L. (2010). Solar wind propagation delay:
660 Comment on "Minimum variance analysis-based propagation of the solar wind observations: Application to
661 real-time global magnetohydrodynamic simulations" by A. Space Weather, 8(6), S06005.
662 https://doi.org/10.1029/2009SW000542

663 Karlsson, T., Trollvik, H., Raptis, S., Nilsson, H., and Madanian, H. (2022). Solar wind magnetic holes can
664 cross the bow shock and enter the magnetosheath, Ann. Geophys., 40, 687–699, https://doi.org/10.5194/angeo-
665 40-687-2022.

666 Julia A. Kropotina, Lee Webster, Anton V. Artemyev, Andrei M. Bykov, Dmitri L. Vainchtein and Ivan Y.
667 Vasko (2021). Solar Wind Discontinuity Transformation at the Bow Shock. The Astrophysical Journal 913 (2)
668 142. DOI: https://dx.doi.org/10.3847/1538-4357/abf6c7 .

669 LeCun, Y., K. Kavukcuoglu, C. Farabet (2010). Convolutional networks and applications in vision, Proceedings
670 of IEEE International Symposium on Circuits and Systems 253–256. DOI: 10.1109/ISCAS.2010.5537907

671 Li, G. (2008). Identifying current-sheet-like structures in the solar wind. The Astrophysical Journal, 672(1),
672 L65–L68. https://doi.org/10.1086/525847

673 Li, H., Wang, C., Cui, T., & Xu, F.(2020). Machine learning approach for solar wind categorization. Earth and
674 Space Science.7, https://doi.org/10.1029/2019EA000997

675 Mailyan, B., Munteanu, C., & Haaland, S. (2008). What is the best method to calculate the solar wind
676 propagation delay? Annals of Geophysics, 26(8), 2383–2394. https://doi.org/10.5194/angeo-26-2383-2008

677 Miao, B., Peng, B., & Li, G. (2011). Current sheets from Ulysses observation. Annals of Geophysics, 29(2),
678 237–249. https://doi.org/10.5194/angeo-29-237-2011

679 Munteanu, C. (2017). Turbulent fluctuations and discontinuities in the solar wind: Statistical properties and
680 possible effects on the terrestrial plasma environment, Ph.D. Thesis, Faculty of Physics, Doctoral School of
681 Physics, University of Bucharest, Romania. Retrieved from
682 https://gpsm.spacescience.ro/gpsm/core/doc/publicatii/teze/Costel_Munteanu_UniBuc_thesis_final_version.pdf

683 Munteanu, C., Haaland, S., Mailyan, B., Echim, M., & Mursula, K. (2013). Propagation delay of solar wind
684 discontinuities: Comparing different methods and evaluating the effect of wavelet denoising. Journal of
685 Geophysical Research: Space Physics, 118(7), 3985–3994. https://doi.org/10.1002/jgra.50429.

686 Munteanu, C., Turicu, D. C., Creț, O., & Echim, M. (2022). Detecting discontinuities from in situ space
687 measurements: Method and FPGA implementation. Earth and Space Science, 9, e2022EA002537. https://doi.
688 org/10.1029/2022EA002537

689 Munteanu, C., Kovács, P., & Echim, M. (2023). An integrated nonlinear analysis (INA) software for space
690 plasma turbulence. Earth and Space Science, 10, e2022EA002692. https://doi. org/10.1029/2022EA002692

691 Newman, R., Vainchtein, D., & Artemyev, A. (2020). Solar wind transient currents: Statistical properties and
692 impact on Earth's magnetosphere. Solar Physics, 295(9), 129. https://doi.org/10.1007/s11207-020-01695-z

693 Paschmann, G., Haaland, S., Sonnerup, B., & Knetter, T. (2013). Discontinuities and Alfvénic fluctuations in
694 the solar wind. Annals of Geophysics, 31(5), 871–887. https://doi.org/10.5194/angeo-31-871-2013

695 Powers, D. M. (2020). Evaluation: from precision, recall and F-measure to ROC, informedness, markedness and
696 correlation. https://doi.org/10.48550/arXiv.2010.16061

697 Tsurutani, B. T., Lakhina, G. S., Verkhoglyadova, O. P., Gonzalez, W. D., Echer, E., & Guarnieri, F. L. (2011).
698 A review of interplanetary discontinuities and their geomagnetic effects. Journal of Atmospheric and Solar-
699 Terrestrial Physics, 73(1), 5–19. https://doi.org/10.1016/j. jastp.2010.04.001

700 Tsurutani, B. T., & Smith, E. J. (1979). Interplanetary discontinuities: Temporal variations and the radial
701 gradient from 1 to 8.5 AU. Journal of Geophysical Research, 84(A6), 2773–2787.
702 https://doi.org/10.1029/JA084iA06p02773

703 Upendran, V., Cheung, M. C. M., Hanasoge, S., & Krishnamurthi, G. (2020). Solar wind prediction using deep
704 learning. Space Weather, 18, e2020SW002478. https://doi.org/10.1029/2020SW002478

705 Vasquez, B. J., Abramenko, V. I., Haggerty, D. K., & Smith, C. W. (2007). Numerous small magnetic field
706 discontinuities of Bartels rotation 2286 and the potential role of Alfvénic turbulence. Journal of Geophysical
707 Research, 112(A11), A11102. https://doi.org/10.1029/2007JA012504

708 Webster, L., Vainchtein, D. & Artemyev, A. (2021). Solar Wind Discontinuity Interaction with the Bow Shock:
709 Current Density Growth and Dawn-Dusk Asymmetry. Sol Phys 296, 87. https://doi.org/10.1007/s11207-021-
710 01824-2

711