



Visualizing and Interpreting Unsupervised Solar Wind Classifications

Jorge Amaya*, Romain Dupuis, Maria Elena Innocenti and Giovanni Lapenta

Mathematics Department, Centre for Mathematical Plasma-Astrophysics, KU Leuven, Leuven, Belgium

OPEN ACCESS

Edited by:

Veronique A. Delouille,
Royal Observatory of Belgium,
Belgium

Reviewed by:

Hui Li,
National Space Science Center (CAS),
China
Xochitl Blanco-Cano,
National Autonomous University of
Mexico, Mexico
Verena Heidrich-Meisner,
University of Kiel, Germany

*Correspondence:

Jorge Amaya
jorge.amaya@kuleuven.be;
jorgeluis.amaya@gmail.com

Specialty section:

This article was submitted to
Space Physics,
a section of the journal
Frontiers in Astronomy and Space
Sciences

Received: 17 April 2020

Accepted: 17 August 2020

Published: 25 September 2020

Citation:

Amaya J, Dupuis R, Innocenti ME and
Lapenta G (2020) Visualizing and
Interpreting Unsupervised Solar Wind
Classifications.
Front. Astron. Space Sci. 7:553207.
doi: 10.3389/fspas.2020.553207

One of the goals of machine learning is to eliminate tedious and arduous repetitive work. The manual and semi-automatic classification of millions of hours of solar wind data from multiple missions can be replaced by automatic algorithms that can discover, in mountains of multi-dimensional data, the real differences in the solar wind properties. In this paper we present how unsupervised clustering techniques can be used to segregate different types of solar wind. We propose the use of advanced data reduction methods to pre-process the data, and we introduce the use of Self-Organizing Maps to visualize and interpret 14 years of ACE data. Finally, we show how these techniques can potentially be used to uncover hidden information, and how they compare with previous empirical categorizations.

Keywords: solar wind, ACE, Self-Organizing Maps, clustering, autoencoder, PCA, unsupervised, machine learning

1. INTRODUCTION

The effects of solar activity on the magnetic environment of the Earth have been observed since the publication of Edward Sabine's work in 1852 (Sabine, 1852). During almost 200 years we have learned about the intimate connection between our star and the plasma environment of the Earth. Three main physical processes connect us to the Sun: the transfer of electromagnetic radiation, the transport of energetic particles, and the flow of solar wind. The later is a continuous stream of charged particles that carries the solar magnetic field out of the corona and into the interplanetary space.

The name *solar wind* was coined by Parker in 1958 because “the gross dynamical properties of the outward streaming gas [from the Sun] are hydrodynamic in character” (Parker, 1958). Over time we have learned that the wind also has many more complex properties. Initially, it was natural to classify the solar wind by defining a boundary between *fast* and *slow* winds (Neugebauer and Snyder, 1966; Schwenn, 1983; Schwenn and Marsch, 1990; Habbal et al., 1997). The former has been associated with mean speed values of 750 km/s (or in some publications with values larger than 600 km/s), while the later shows a limit at 500 km/s, where the compositional ratio (Fe/O) shows a break (Feldman et al., 2005; Stakhiv et al., 2015). The solar wind also carries information about its origins on the Sun. At certain solar distances the ion composition of the solar wind is expected to be frozen-in, reflecting the electron temperature in the corona and its region of origin (Feldman et al., 2005; Zhao et al., 2009; Stakhiv et al., 2015). These particles have multiple energies and show a variety of kinetic properties, including non-Maxwellian velocity distributions (Pierrard and Lazar, 2010; Matteini et al., 2012).

The solar wind is also connected to the Sun by Interplanetary Magnetic Field (IMF) lines directed toward the Sun, away from the Sun, or in the case of flux ropes, connected at both ends (Gosling et al., 2010; Owens, 2016). The region separating IMF lines of opposite polarity (directed away or toward the Sun) is called the Heliospheric Current Sheet (HCS) (Smith, 2001). When a spacecraft crosses the HCS instruments onboard measure the change in polarity of the magnetic field. In quiet wind conditions the plasma around the HCS presents discontinuities in density, temperature velocity and magnetic field (Eselevich and Filippov, 1988). This perturbed region surrounding the HCS is called the Heliospheric Plasma Sheet (HPS). The passage of the spacecraft from one side of the HPS to the other is known as a Sector Boundary Crossing (SBC) (Winterhalter et al., 1994). In spacecraft observations these are sometimes confused with Corotating Interaction Regions (CIR), which are zones of the solar wind where fast flows have caught up with slow downstream solar wind, compressing the plasma (Fisk and Lee, 1980; Richardson, 2004).

From the point of view of a spacecraft SBCs and CIRs can show similar sudden changes in the plasma properties. These two in turn are often grouped and mixed with other transient events, like Coronal Mass Ejections (CME) and Magnetic Clouds (MC). Since 1981 when Burlaga et al. (1981) described the propagation of MC behind an interplanetary shock, it was suspected that CMEs and MC were coupled. However, more recent studies show that CMEs observed near the Sun do not necessarily become MC, but instead “pressure pulses” (Gopalswamy et al., 1998; Wu et al., 2006).

Much more recently it has been revealed, by observations from Parker Solar Probe, that the properties of the solar wind can be drastically different closer to the Sun, where the plasma flow is more pristine and has not yet mixed with the interplanetary environment. Patches of large intermittent magnetic field reversals, associated with jets of plasma and enhanced Poynting flux, have been observed and named “switchbacks” (Bale et al., 2019; Bandyopadhyay et al., 2020).

The solar wind is thus not only an hydrodynamic flow, but a compressible mix of different populations of charged particles and electromagnetic fields that carry information of their solar origin (helmet streamer, coronal holes, filaments, solar active regions, etc.) and is the dominion of complex plasma interactions (ICMEs, MC, CIRs, SBCs, switchbacks).

To identify and study each one of these phenomena we have relied in the past on a manual search, identification and classification of spacecraft data. Multiple authors have created empirical methods of wind type identification based on *in-situ* satellite observations and remote imaging of the solar corona. Over the years the number and types of solar wind classes has changed, following our understanding of the complexity of heliospheric physics.

Solar wind classification serves four main roles:

1. it is used for the characterization of its origins in the corona,
2. to identify the conditions where the solar wind is geoeffective,
3. to isolate different plasma populations in order to perform statistical analysis,

4. to study the basic transport effects of space plasmas of different nature.

Among the existing classifications we can include the original review work by Withbroe (1986), the impressive continuous inventory by Richardson et al. (2000) and Richardson and Cane (2010, 2012), and the detailed studies by Zhao et al. (2009) and (Xu and Borovsky, 2015). These publications classify the solar wind based on their ion composition, and on the transient events detected. Each system includes two, three, or four classes, generally involving coronal-hole origins, CMEs, streamer belt origins, and sector reversal regions.

The precise point of origin of the solar wind can be traced back from spacecraft positions to the solar corona and the photosphere: multiple authors (Neugebauer et al., 2002; Zhao et al., 2009, 2017; Fu et al., 2015) have used a ballistic approximation coupled to a Potential Field Source Surface (PFSS) model to trace back solar wind observations to their original sources on the Sun. This procedure relies on multiple assumptions, including a constant solar wind speed and a force free magnetic field configuration of the solar corona. The uncertainty on the source position is estimated around $\pm 10^\circ$ by Neugebauer et al. (2002). This is currently the best method to acquire the *ground truth* about the origin of the solar wind. Unfortunately, to our knowledge, there is no central repository of solar wind origins for any space mission that we can use to train or verify our novel machine learning techniques.

We are moving now toward a new era of data analysis, where manual human intervention can be replaced by *intelligent* software. The trend has already started, with the work by Camporeale et al. (2017) who used (Xu and Borovsky, 2015) classes to train a Gaussian Process algorithm that autonomously assigns the solar wind to the proper class, and by Roberts et al. (2020) who used unsupervised classification to perform a 4 and 8 class solar wind classification. A recent publication by Bloch et al. (2020) uses unsupervised techniques to classify ACE and Ulysses observations, and Li et al. (2020) have successfully tested 10 different supervised techniques to reproduce the categories introduced by Xu and Borovsky (2015).

The most basic ML techniques learn using two approaches: (A) in supervised learning the algorithms are shown a group of inputs, $\mathbf{X} \in \mathbb{R}^n$, and outputs, $\mathbf{Y} \in \mathbb{R}^o$, with the goal of finding a non-linear relationship between them, $\xi_s: \mathbf{X} \rightarrow \mathbf{Y}$, (B) in unsupervised learning the machine is presented with a cloud of multi-dimensional points, $\mathbf{X} \in \mathbb{R}^n$, that have to be autonomously categorized in different classes, either performing associations with representative points in the same data space, $\xi_u: \mathbf{X} \rightarrow \mathbf{W} \in \mathbb{R}^n$, or by grouping neighboring data points together into an assigned set, $\xi_u: \mathbf{X} \rightarrow g \in \mathbb{R}$. This means that we can program the computer to learn about the different types of solar wind using the existing empirical classifications using method (a), or allowing the computer to independently detect patterns in the solar wind properties with method (b).

In the present work we show how the second method, unsupervised classification, can be used to segregate different types of solar wind. In addition, we show how to visualize and interpret such results. The goal of this paper is to introduce the

use of unsupervised techniques to our community, including the best use practices and the opportunities that such methods can bring. We promote the use of one specific type of classification, called Self-Organizing Maps, and we compare it to simpler classification techniques.

In the next sections we present in detail the techniques of data processing (section 2.1), data dimension reduction (sections 2.2.1, 2.2.2 and 2.2.3) and data clustering (section 2.2.4) that we have used. We then present in detail the Self-Organizing Map technique and all its properties in section 2.2.5. We show how to connect all of these parts together in section 2.2.6, and finally we show how the full system can be used to study 14 years of solar wind data from the ACE spacecraft in section 3.

2. MATERIALS AND METHODS

2.1. Data and Processing

2.1.1. Data Set Used

The solar wind data used in this work was obtained by the Advanced Composition Explorer (ACE) spacecraft, during a period of 14 years, between 1998 and 2011. The data can be downloaded from the FTP servers of The ACE Science Center (ASC) (Garrard et al., 1998). The files in this repository correspond to a compilation of hourly average data from four instruments: MAG (Magnetometer) (Smith et al., 1998), SWEPAM (Solar Wind Electron, Proton, and Alpha Monitor) (McComas et al., 1998), EPAM (Electron, Proton, and Alpha Monitor) (Gold et al., 1998), and SWICS (Solar Wind Ion Composition Spectrometer) (Gloeckler et al., 1998). A detailed description of the entries in this data set can be found in the ASC website listed in the Data Availability Statement.

A total of 122,712 data points are available. However, routine maintenance operations, low statistics, instrument saturation and degradation produce gaps and errors in the data. The SWICS data includes a flag assessing the quality of the calculated plasma moments. We retain only *Good quality* entries. Our pre-processed data set contains a total of 72,454 points.

2.1.2. Additional Derived Features

We created additional features for each entry, based on previous knowledge of the physical properties of the solar wind. Some are derived from the existing properties in the data set, others computed from statistical analysis of their evolution. We introduce here the additional *engineered* features included in our data set.

Multiple techniques have been proposed in the literature to identify ejecta, Interplanetary Coronal Mass Ejections (ICME), and solar wind origins in the ACE data. Zhao et al. (2009) suggest that, during solar cycle 23, three classes of solar wind can be identified using its speed, V_{sw} , and the oxygen ion charge state ratio, O^{7+}/O^{6+} . It has been shown that slow winds originating in coronal streamers correlate with high values of the charge state ratio and fast winds coming from coronal holes present low values (Schwenn, 1983; Withbroe, 1986; Schwenn and Marsch, 1990). Plasma formed in coronal loops associated with CMEs also show high values of the charge state ratio (Xu and Borovsky, 2015;

Zhao et al., 2017). The classification boundaries of the Z09 model, proposed by Zhao et al. (2009), are presented in **Table 1**.

Xu and Borovsky (2015) suggested an alternative four classes system based on the proton-specific entropy, $S_p = T_p/n_p^{2/3}$ [K cm²], the Alfvén speed, $V_A = B/(\mu_0 m_p n_p)^{1/2}$ [Km s⁻¹], and the ratio between the expected and the measured proton temperature, $T_{exp}/T_p = (V_{sw}/258)^{3.113}/T_p$ [-], where n_p is the proton number density, m_p is the proton mass, and μ_0 is the permeability of free space. The classification boundaries used for the X15 model, proposed by Xu and Borovsky (2015), are also presented in **Table 1**. For each entry in the data set we have included the values of S_p , V_A , T_{exp} , $T_{ratio} = T_{exp}/T_p$, and the solar wind type.

Two additional empirical threshold methods will be included in this work for comparison. These two methods were derived from the compositional observations of the solar wind at higher heliospheric latitudes, using data from the *Ulysses* mission (Wenzel et al., 1992). The first model, that we call vS15, comes from the work by von Steiger and Zurbuchen (2015), where the first figure shows a clear division between Coronal Hole (CH) sources and non-Coronal Hole (NCH) wind. The boundary between the two classes is presented in **Table 1**. The second threshold model was presented as an example by Bloch et al. (2020). This boundary, named here B20, is an empirical approximation that divides CH and NCH origin winds. The threshold values are shown in **Table 1**.

In addition to the instantaneous properties of the solar wind used in all previous classifications, we can perform statistical operations over a window of time of 6 h, including values of the maximum, minimum, mean, standard deviation, variance, auto-correlation, and range. We expect to capture with some of these quantities turbulent signals or sudden jumps associated with different transient events. These additional rolling operations are a complement to the stationary solar wind parameters mentioned above and add information about the temporal evolution of the plasma. The selection of the statistical parameters and the window of time is arbitrary and will require a closer examination in the future.

An additional term, which has been successfully used in the study of solar wind turbulence (D'Amicis and Bruno, 2015; Zhao et al., 2018; Magyar et al., 2019; Adhikari et al., 2020), is included here to account for additional time correlations. The normalized cross-helicity, σ_c , is defined in Equation (1), where $\mathbf{b} = (\mathbf{B} - \langle \mathbf{B} \rangle) / (\mu_0 m_p n_p)^{1/2}$ is the fluctuating magnetic field in Alfvén units, $\mathbf{v} = \mathbf{V}_{sw} - \langle \mathbf{V}_{sw} \rangle$ is the fluctuating solar wind velocity, and $\langle \cdot \rangle$ denotes the averaging of quantities over a time window of 3 h (Roberts et al., 2020).

$$\sigma_c = 2 (\mathbf{b} \cdot \mathbf{v}) / (\mathbf{b}^2 + \mathbf{v}^2) \quad (1)$$

Due to gaps in the data, some of the above quantities can not be obtained. We eliminate from the data set all entries for which the derived features presented in this section could not be calculated. This leaves a total of 51,374 entries in the data set used in the present work.

TABLE 1 | Solar wind types and boundaries as defined by the empirical models: Z09, X15, vS15, and B20.

#	SW type	Condition	References
0	CH	$\log_{10} O^{7+}/O^{6+} \leq 0.145$	Zhao et al., 2009
2	NCH	$0.145 < \log_{10} O^{7+}/O^{6+} < 6.008e^{(-0.00578V_{sw})}$	
4	TR	$\log_{10} O^{7+}/O^{6+} > 6.008e^{(-0.00578V_{sw})}$	
0	CH	Not type TR, and $\log_{10}(S_p) > -0.525 \log_{10}(T_{exp}/T_p) - 0.676 \log_{10}(V_A) + 1.74$	Xu and Borovsky, 2015
2	SB	Not type CH, TR, or SR	
3	SR	Not type TR, and $\log_{10}(S_p) < -0.125 \log_{10}(T_{exp}/T_p) - 0.658 \log_{10}(V_A) + 1.04$	
4	TR	$\log_{10}(V_A) > 0.055 \log_{10}(T_{exp}/T_p) + 0.277 \log_{10}(S_p) + 1.83$	
0	CH	$\log_{10} O^{7+}/O^{6+} \times \log_{10} C^{6+}/C^{5+} \leq 0.01$	von Steiger and Zurbuchen, 2015
2	NCH	Not type CH	
0	CH	$1.25 \times \log_{10} O^{7+}/O^{6+} + 6.75 < \log_{10}(S_p)$	Bloch et al., 2020
2	NCH	Not type CH	

The four types are: fast solar wind of coronal hole origin (CH), slow wind of non-coronal hole origin (NCH), transients, including ejecta, ICMEs, CIRs, MCs or other sudden jumps in solar wind parameters (TR), solar wind originated in the streamer belt (SB), and solar wind of sector reversal origins (SR). The ID value in the first column is arbitrary and has been chosen to simplify the visualization of our results.

To account for the differences in units and scale, each feature column F in the data set is normalized to values between 0 and 1, using: $f = (F - \min F) / (\max F - \min F)$.

Not all the features might be useful and some of them can be strongly correlated. We do not perform here a detailed evaluation of the inter dependencies of the different features, and we leave that task for a future work. The present manuscript focuses on the description of the methodology and on the visualization and interpretation capabilities of unsupervised machine learning classification. We limit our work here to test and compare a single model that incorporates a total of 15 features. These are listed in Table 2.

2.1.3. Complementary Data Catalogs

We support the interpretation of our results using data from three solar wind event catalogs. The first is the well-known Cane and Richardson catalog that contains information about ICMEs detected in the solar wind in front of the Earth (Cane and Richardson, 2003; Richardson and Cane, 2010)¹. We used the August 16, 2019 revision. As the authors state in their website, there is no spreadsheet or text version of this catalog and offline editing was necessary. We downloaded and re-formatted the catalog to use it in our application. The CSV file created has been made available in our repository. We call this, the Richardson and Cane catalog.

The second catalog corresponds to the ACE List of Disturbances and Transients² produced by the University of

New Hampshire. As in the previous case, the catalog is only available as an html webpage, so we have manually edited the file and extracted the catalog data into a file also available in our repository. This is hereafter referred to as the UNH catalog.

Finally, we also included data from the Shock Database³ maintained by Dr. Michael L. Stevens and Professor Justin C. Kasper at the Harvard-Smithsonian Center for Astrophysics. Once again we have gathered and edited multiple web-pages in a single file available in our repository. In this work this database will be known as the CfA catalog.

2.2. Dimension Reduction and Clustering

2.2.1. Dimension Reduction Using PCA

Principal Component Analysis (PCA) is a mathematical tool used in data analysis to simplify and extract the most relevant features in a complex data set. This technique is used to create entries composed of linearly independent *principal components*. These are the eigenvectors, \mathbf{v} , of the covariance matrix $\Sigma = (\Sigma_{ij})$ applied to the centered data (Equation 2), ordered from the largest to the smallest eigenvalue, $\lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_n$, where \bar{X} is the mean value of each one of the n original features (Equation 3), and m is the total number of entries in the data set. The projection of the data onto the principal component space ensures a maximal variance on the direction of the first component. Each subsequent principal component is orthogonal to the previous ones and points in the direction of maximal

¹Near-Earth Interplanetary Coronal Mass Ejections Since January 1996: <http://www.srl.caltech.edu/ACE/ASC/DATA/level3/icmetable 2.html>

²ACE Lists of Disturbances and Transients: http://www.ssg.sr.unh.edu/mag/ace/ACElists/obs_list.html

³Harvard-Smithsonian, Center for Astrophysics, Interplanetary Shock Database—ACE: https://www.cfa.harvard.edu/shocks/ac_master_data/

TABLE 2 | List of features used for the AE+DSOM model.

ID	Name in the database	AE + DSOM
0	proton_speed	✓
1	proton_density	✓
2	O7to6	✓
3	C6to5	✓
4	FetoO	✓
5	avqFe	✓
6	proton_temp	✓
7	sigmac ^(*)	✓
8	Sp	✓
9	Va	✓
10	Tratio	✓
11	proton_speed_range	✓
12	Bn_range	✓
13	FetoO_range	✓
14	O7to6_range	✓
	Initial year	1998
	Final year	2011
	Neurons / encoding layer	[15, 10]
	Optimizer	Adam
	Learning rate	0.042
	Lattice nodes	10×10
	ε	0.005
	η	3.0

The logarithm of all quantities was used, except for the features marked with an asterisk (*). Bottom: data range and hyper-parameters of the AE and the SOM.

variance in the residual sub-space (Shlens et al., 2014).

$$\Sigma_{ij} = \frac{1}{m} \sum_{k=1}^m (\mathbf{X}_i^k - \bar{\mathbf{X}}_i) (\mathbf{X}_j^k - \bar{\mathbf{X}}_j) \quad (2)$$

$$\bar{\mathbf{X}} = \frac{1}{m} \sum_{i=1}^m \mathbf{X}_i \quad (3)$$

$$\Sigma \mathbf{v} = \lambda \mathbf{v} \quad (4)$$

The PCA transformation creates the same number of components in the transformed space, $\tilde{\mathbf{X}}$, as features in the original data space \mathbf{X} . However, components with small eigenvalues belong to a dimension where the variance is so small that it is impossible to separate points in the data. It is a general practice in data reduction to keep only the first k components that explain at least a significant portion of the total variance of the data, $\sum_{i=1}^k \lambda_i / \text{Tr}(\Sigma) > \epsilon$. This allows for a selection of information that will effectively differentiate data points, and for a reduction of the amount of data to process during analysis. Many techniques have been suggested for the selection of the values of k and the cut-off ϵ (Rea and Rea, 2016). We use the value of $\epsilon = 0.95$.

2.2.2. Dimension Reduction Using Kernel PCA

PCA has a limitation: the principal components are a linear combination of the original properties of the solar wind. The Kernel PCA (KPCA) is an extension of the PCA that allows to perform non-linear transformations of the original data. The goal in KPCA is to perform the original PCA operations in a high dimensional space.

For a list of m data points composed of n features, it is sometimes difficult (or impossible) to build a linear hyper-plane that dissects regions of different density. However, it is possible to conceive a function, $\xi: \mathbf{X} \in \mathbb{R}^n \rightarrow \tilde{\mathbf{X}} \in \mathbb{R}^m$, that will transform all the data into a space where each cluster of points can be linearly separable. The goal is then to avoid explicitly calculating the high-dimensional function ξ by building a Kernel, \mathbf{K} , which is the inner product of the high-dimensional space:

$$\mathbf{K} = k(\mathbf{X}_i, \mathbf{X}_j) = \xi(\mathbf{X}_i)^T \xi(\mathbf{X}_j) \quad (5)$$

In this space the projected points are linearly separable using the same principles of the PCA. In this case the covariance matrix would be expressed as:

$$\Sigma_{ij} = \frac{1}{m} \sum_{k=1}^m \xi(\mathbf{X}_i) \xi(\mathbf{X}_j)^T \quad (6)$$

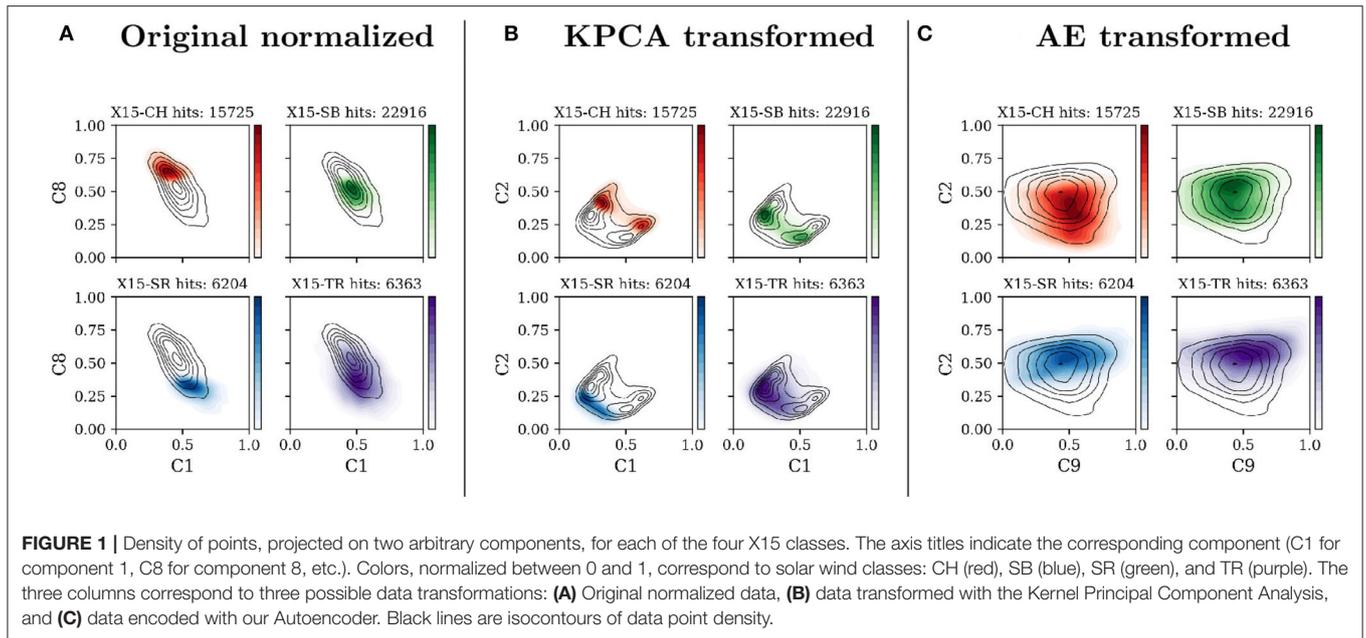
$$\mathbf{v} = \sum_{i=1}^n a_i \xi(x_i) \quad (7)$$

Popular kernel functions include Gaussian, polynomial and hyperbolic tangent. The transformation is reduced to solving the eigenvalue problem: $\mathbf{K}\mathbf{a} = \lambda\mathbf{a}$, where \mathbf{a} are the coefficients of the linear combination of the eigenvectors (Equation 7). Although a powerful tool, KPCA requires the creation of an $m \times m$ matrix that can consume large amounts of time and memory resources.

In this work we use KPCA with a polynomial kernel of order eight (8). We also apply the procedure described before to select the total number of retained components: we impose $\epsilon = 0.95$. Cutting off the number of components implies a loss of data. To verify that only minimal information is lost, we perform a transformation of all our data set followed by an inverse transformation. The relative error between the two is normally distributed around zero with <1% of variance.

2.2.3. Dimension Reduction Using Autoencoders

An alternative to data reduction is the use of Autoencoders (AE). These are machine learning techniques that can create non-linear combinations of the original features projected on a latent space with less dimensions (Hinton and Salakhutdinov, 2006). This is accomplished by creating a system where an encoding function, ϕ , maps the original data \mathbf{X} to a latent space, $\mathcal{F} \in \mathbb{R}^d$ (Equation 8). A decoder function, ψ , then maps the latent space back to the original input space (Equation 9). The objective of the autoencoder is to minimize the error between the original data and the data produced by the compression-decompression



procedure as shown in Equation (10).

$$\phi : \mathbf{X} \rightarrow \mathbf{Z} \in \mathcal{F} \quad (8)$$

$$\psi : \mathbf{Z} \in \mathcal{F} \rightarrow \mathbf{X} \quad (9)$$

$$\phi, \psi = \arg \min_{\phi, \psi} \|\mathbf{X} - (\phi \circ \psi)\mathbf{X}\|^2 \quad (10)$$

Autoencoders can be represented as feed-forward neural networks, where fully connected layers lead to a central bottleneck layer with few nodes and then expands to reach again the input layer size. An encoded element, $z \in \mathcal{F}$, can be obtained from a data entry, $x \in \mathbf{X}$, following the standard neural network function (Equation 11), where \mathbf{W} is the weights matrix, c is the bias, and f is the non-linear activation function.

$$z = f(\mathbf{W}x + c) \quad (11)$$

$$\hat{x} = f'(\mathbf{W}'z + c') \quad (12)$$

$$\mathcal{L}(x, \hat{x}) = \|x - \hat{x}\|^2 \quad (13)$$

The decoding procedure, shown in Equation (12), transforms $z \rightarrow \hat{x}$, where the prime quantities are associated with the decoder. The loss function, $\mathcal{L}(x, \hat{x})$, is the objective to be minimized by the training of the neural network using gradient descent. Once training is completed, the vector z is a projection of the input vector x onto the lower dimensional space \mathcal{F} .

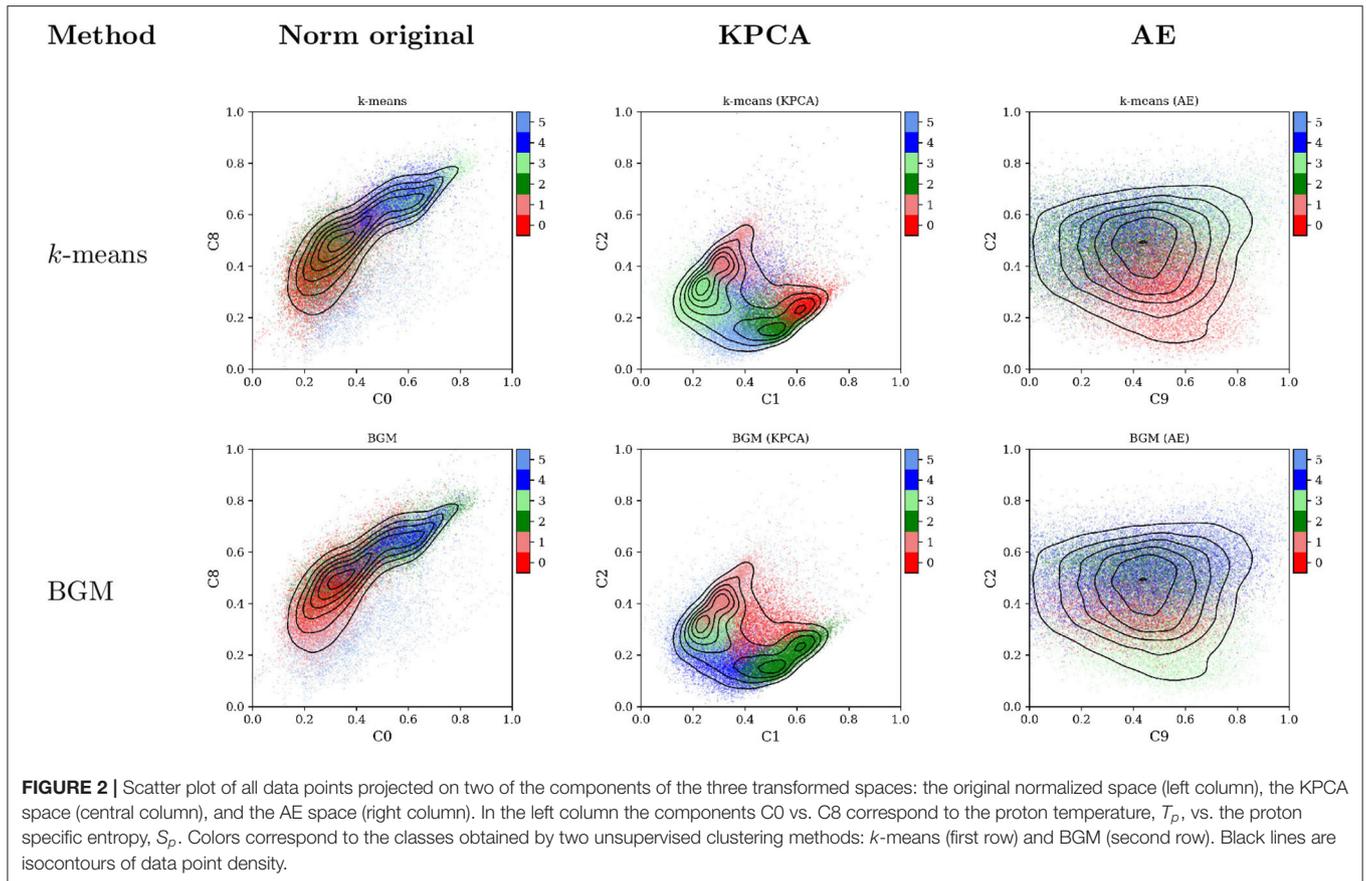
Additional enhancements and variations of this simple autoencoder setup exist in the literature, including multiple regularization techniques to minimize over-fitting (Liang and Liu, 2015), Variational Autoencoders (VAE) that produce encoded Gaussian distribution functions (Kingma and Welling, 2013), and Generative Adversarial Networks that automatically generate new data (Goodfellow et al., 2014). In this work we use the most basic form of autoencoders, presented above.

In the present work we will be showing different representations of the solar wind data, transformed with different techniques and projected on flat planes. **Figure 1** presents our data set in three different projections: (A) the original feature space, normalized between zero and one, (B) the transformed data set using the KPCA method, and (C) the AE transformed data. In each panel four histograms present the distribution of the X15 classes, on two arbitrary components identified by the axis title.

2.2.4. Clustering Techniques

The goal of unsupervised machine learning is to group data points in a limited number of clusters in the N -dimensional space $\Omega \in \mathbb{R}^n$, where n is the number of features (components or properties) in the data set. Multiple techniques can be used to perform multi-dimensional clustering. We present in **Figure 2** the application of two basic clustering techniques to classify our data set. Following the same order as before, the first column in the figure contains all data points projected in the original normalized feature space; column two contains scatter plots of the points after KPCA transformation; column three contains the same points encoded in the AE latent space. Each row corresponds to a different clustering method. The colors in the top row were obtained using the k -means method (Lloyd, 1982), while the colors in bottom panels were obtained using the Bayesian Gaussian Mixture (BGM) (Bishop, 2006).

The k -means technique has already been used in multiple publications for the determination of solar wind states (Heidrich-Meisner and Wimmer-Schweingruber, 2018; Roberts et al., 2020). The BGM technique has also been recently used by Bloch et al. (2020) to classify solar wind observations by the ACE and Ulysses missions. Mixture models similar to the BGM have also been recently used to classify space plasma regions in magnetic reconnection zones (Dupuis et al., 2020). None of these previous



publications used data transformation to solve the classification problem in a more suitable latent space.

The colors used in **Figure 2** are assigned randomly by each clustering technique. The most glaring issue with them is that different methods can lead to different clusters of points. The BGM and the *k*-means do not agree on their classification in the PCA and the AE space. More importantly, for each technique, slight modifications of the clustering parameters, e.g., using a different seed for the random number generator, can lead to very different results. We address this last issue using an algorithm that launches the *k*-means (the BGM) algorithm 100 (30) times until the method converges to a global minimum. The final results are implementation dependent.

In the present data set, the cloud of points is convex and well-distributed in all components. This raises one additional issue, observed more clearly in the second column of **Figure 2**: when classical clustering methods are applied to relatively homogeneously dense data, it divides the feature space in Voronoï regions with linear hyper-plane boundaries. This is an issue with all clustering techniques based on discrimination of groups using their relative distances (to a centroid or to the mean of the distribution). To avoid this problem density-based techniques, such as DBSCAN (Ester et al., 1996), and agglomeration clustering methods, use a different approach.

However, we can not apply them here because in such homogeneous cloud of points these techniques lead to a trivial solution where all data points are assigned to a single class. An alternative projection was used by Bloch et al. (2020), who performed a Uniform Manifold Approximation and Projection (UMAP). We performed the same projection unsuccessfully: the Ulysses data used in that publication contains a very dense and large number of CH observations. ACE lacks such a rich variety of CH data, so applying a UMAP leads to a single class.

There is no guarantee that a single classification method, with a particular set of parameters will converge to a physically meaningful classification of the data if the points in the data do not have some level of separability, or have multiple zones of high density. This is also true for other classification methods based on *supervised learning*. The same issues will be observed when the training data include target classes derived from dense data clouds using simple hyper-plane boundaries, as done for the Z09 and X15 classes. An example of such application was published by Camporeale et al. (2017) and Li et al. (2020). The authors used the X15 classification to train supervised classifiers. No new information is gained with such methods, as the empirical boundaries are already mathematically known. A more compelling task would be to compare all classification methods against a *ground truth*, i.e., against a catalog of footprint

locations on the solar surface. But such catalog, to our knowledge, does not exist.

2.2.5. Self-Organizing Maps

2.2.5.1. Classical SOM

Following the definitions and notations by Villmann and Claussen (2006), a class can be defined as $C_i \stackrel{\text{def}}{=} \{x \in \Omega | \Phi(x) = w_i\}$, where Φ is a function from Ω to a finite subset of k points $\{w_i \in \mathbb{R}^N\}_{i=1..k}$. A cluster C_i is then a partition of Ω , and $\{w_i\}$ are the code words (also known as nodes, weights or centroids) associated. The mapping from the data space to the code word set, $\Phi : \Omega \rightarrow \mathcal{W}$, is obtained by finding the closest neighbor between the points x and the code words w (Equation 14). The code word w_s , the closest node to the input x_s , is called the *winning element*. The class C_i corresponds to a Voronoï region of Ω with center in w_i .

$$\Phi : x \rightarrow \arg \min_{i \in \mathcal{N}} (\|x - w_i\|) \quad (14)$$

A Self-Organizing Map (SOM) is also composed of structured nodes arranged in a lattice, each one assigned to a fixed position p_i in \mathbb{R}^q , where q is the dimension of the lattice (generally $q = 2$). The map nodes are characterized by their associated code words. The SOM learns by adjusting the code words w_i as input data x is presented.

The SOM is the ensemble of code words and nodes $\{w_i, p_i\} \in (\Omega \times \mathbb{R}^q)$. For a particular entry x_s , the code word $s \in \mathcal{N}$ is associated to the winning node p_s , if the closest word to x_s is w_s . At every iteration of the method, all code words of the SOM are shifted toward x following the rule:

$$\Delta w_i = \epsilon(t) h_\sigma(t, i, s)(x - w_i) \quad (15)$$

with $h_\sigma(t, i, j)$ defined as the lattice neighbor function:

$$h_\sigma(t, i, j) = e^{-\frac{\|p_i - p_j\|^2}{2\sigma(t)^2}} \quad (16)$$

where $\epsilon(t)$ is the time dependent learning rate (Equation 17), and $\sigma(t)$ is the time dependent lattice neighbor width (Equation 18). The training of the SOM is an iterative process where each data point in the data set is presented to the algorithm multiple times $t = 0, 1, \dots, t_f$. In these equations the subscript 0 refers to initial values at $t = 0$ and the subscript f to values at $t = t_f$.

$$\epsilon(t) = \epsilon_0 \left(\frac{\epsilon_f}{\epsilon_0} \right)^{t/t_f} \quad (17)$$

$$\sigma(t) = \sigma_0 \left(\frac{\sigma_f}{\sigma_0} \right)^{t/t_f} \quad (18)$$

This procedure places the code words in the data space Ω in such a way that neighboring nodes in the lattice are also neighbors in the data space. The lattice can be presented as a q -dimensional image, called map, where nodes sharing similar properties are organized in close proximity.

The main metric for the evaluation of the SOM performance is called the quantization error:

$$Q_E = \frac{1}{m} \sum_{i=1}^m \|x_i - w_{x_i}\| \quad (19)$$

where m , is the total number of entries in the data set. It has been shown that the SOM tends to converge in the mean-square (m.s.) sense to the probabilistic density center of the multi-dimensional input subset (Yin and Allinson, 1995). This means that, if the SOM hyper-parameters are chosen correctly, the code words of the SOM will have a tendency to move toward high density regions of subsets of the input data, and will be located close to the mean of the subset points.

Once the training of the SOM is finished, the code words w_i can be grouped together using any clustering technique, e.g., k -means. The nodes of the SOM with close properties will be made part of the same class. The classes created are an ensemble of Voronoï subspaces, allowing a complex non-linear partitioning of the data space Ω .

The final number of clusters is an input of the algorithm, but can also be calculated autonomously. The Within Cluster Sum of Squares (WCSS) can be used as a metric of the compactness of the clustered nodes. As its name implies the WCSS is the sum of the squared distances from each node to their cluster point. If only one class is selected, the large spread of the nodes would produce a high WCSS. The lowest possible value of the WCSS is obtained for a very high number of classes, when the number of classes is equal to the number of nodes. But such extreme solution is also unpractical. The optimal number of clusters can be obtained using the Kneedle class number determination (Satopaa et al., 2011). We use this automatic technique to let the machine select the optimal number of solar wind classes.

2.2.5.2. Dynamic SOM

The time dependence of the SOM training allows the code words w_i to reach steady coordinates by slowing down their movement over the iterations. Due to the minimization of the distance in Equation (14) code words tend to agglomerate around high density zones of the feature space. The Dynamic Self-Organizing Map (DSOM), introduced by Rougier and Boniface (2011), eliminates the time dependence and allows to cover larger zones of the space outside of the high density regions.

The DSOM is a variation of the SOM where the learning function (Equation 15) and the neighbor function (Equation 16) are replaced by Equations (20) and (21), respectively:

$$\Delta w_i = \epsilon \|x - w_i\|_\Omega h_\eta(i, s, x)(x - w_i) \quad (20)$$

$$h_\eta(i, s, x) = e^{-\frac{1}{\eta^2} \frac{\|p_i - p_j\|^2}{\|x - w_s\|_\Omega^2}} \quad (21)$$

where ϵ is a constant learning rate, $h_\eta(i, s, x)$ is defined as the new lattice neighbor function, and η is the *elasticity* parameter. In their work (Rougier and Boniface, 2011) show that DSOM can be used to draw a larger sample of the feature space Ω , reducing the agglomeration of code words around high density zones. The

main parameters of the DSOM, η and ϵ , control the convergence of the method. A large ϵ moves the code words, \mathbf{w} , very fast with each new iteration; a very low value moves the points slowly in the space. A high elasticity, η , keeps all the nodes extremely close to each other, while a low value does not induce movement on far away code words. The best compromise is to use a very low value of the learning rate coupled with a mid-range elasticity, and a large number of training epochs. This can ensure a relative good convergence to a steady set of code words.

One special advantage of the DSOM is that it can be trained *online*, i.e., it is not necessary to re-train all the model when new data arrives: it adapts automatically to new information.

2.2.5.3. Visualization of SOM and DSOM

Most clustering techniques do not guarantee to converge to a steady immutable solution. Differences in the training parameters or slight changes in the data can have an important impact on the final classification. Clustering tools can be used for statistical analysis, comparisons, data visualization and training of supervised methods. But it will be practically impossible to claim the existence of a general objective set of states discovered only by the use of these basic clustering techniques.

However, SOMs and DSOMs provide an important tool for the study of the solar wind: the maps are composed of nodes that share similar properties with its immediate neighbors. This allows for visual identification of patterns and targeted statistical analysis.

We used the python package MiniSom (Vettigli, 2013) as the starting point of our developments. Multiple methods of the MiniSom have been overloaded to implement the DSOM, and to use a lattice of hexagonal nodes. All auxiliary procedures used to calculate inter-nodal distances, node clustering, data-to-node mapping, and class boundary detection have been implemented by us. All visualization routines are original and have been developed using the python library Matplotlib (Hunter, 2007).

Figure 3 shows the basic types of plots that can be generated using the SOM/DSOM techniques. We present in this figure the outcome of our model, combining a non-linear AE transformation of the ACE data set with the unsupervised classification of the encoded data using the DSOM method. Panel (A) shows a histogram of two components of the feature space Ω , with dots marking the position of the code words \mathbf{w}_i . The colors of the dots represent their DSOM classification. The red lines connect a single code word \mathbf{w}_s with its six closest neighbors. Panel (B) shows the same information as in the previous panel, but using a scatter plot colored by the DSOM classification. This image shows the domain of influence of each one of the DSOM classes.

Panel (C) shows the *hit map* of the DSOM. It contains the lattice nodes \mathbf{p}_i associated to the code words \mathbf{w}_i . They are depicted as hexagons with sizes representing the number of data points connected to each node and colored by their DSOM class. The thickness of the lines between lattice nodes represent the relative distance to its neighbors in the feature space Ω . Red lines connect the node \mathbf{p}_s , associated to the code word \mathbf{w}_s in panel (A), to its closest neighbors.

Figure 3D displays three components of the code words \mathbf{w}_i associated to each one of the \mathbf{p}_i nodes. The node components have been mapped to the basic colors Red, Green and Blue (RGB) and combined together to produce the composite color shown in the figure.

These four representations are only a few examples of the variety of data that can be represented using SOMs. The most important aspect of the SOMs is that data is represented in simple 2D lattices where the nodes share properties with their neighbors. Here we also decided to use hexagonal nodes, connecting 6 equidistant nodes, but other types of representations are also valid, e.g., square or triangular nodes.

2.2.6. The Full Architecture

The previous sections introduced all the individual pieces that we use for the present work. Here, we give a global view of the full model. **Figure 4** shows how all the components are interconnected. At the center of the image is the processed and normalized original ACE data set. The blue dashed lines show the unsupervised techniques already presented by Heidrich-Meisner and Wimmer-Schweingruber (2018), Bloch et al. (2020), and Roberts et al. (2020). The KPCA step is added to the data pipelines used in the literature in order to project the data into a hyper-space where the class boundaries are better defined.

On the right side of the same figure we present our main approach: we perform first a data encoding using an AE, then we perform unsupervised classification of the solar wind with the *k*-means, BGM and DSOM methods. After training, the code words of the DSOM are clustered to group together nodes that share similar properties. This second level classification is done using the *k*-means++ algorithm with 100 re-initializations (it is in general recommended to use between 50 and 500 initializations, searching for a global optimum, as different random runs can lead only to a local minima). We use the Kneede method to automatically select the number of classes that the DSOM will produce (Satopaa et al., 2011). The BGM and the *k*-means clustering techniques are included for comparison.

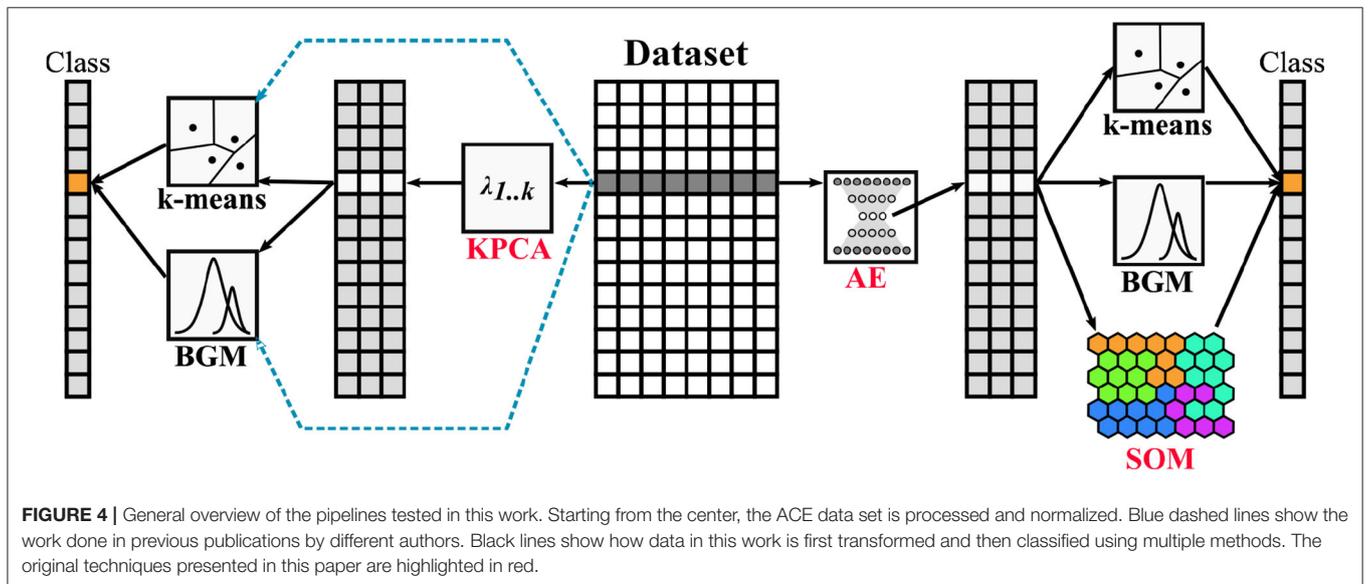
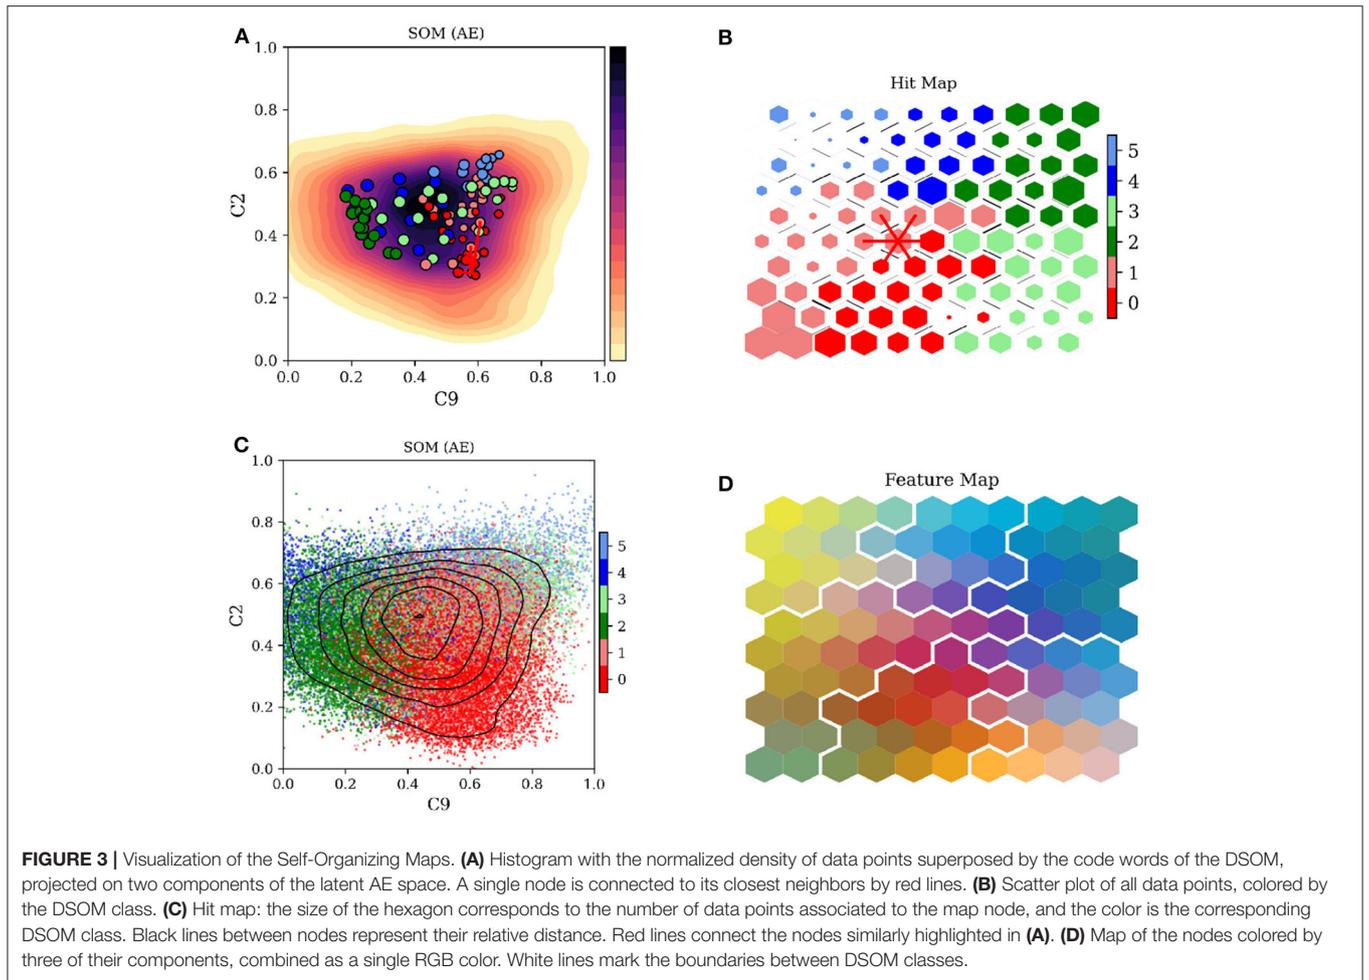
All the software was implemented in Python using as main libraries PyTorch (Paszke et al., 2019), Scikit-learn (Pedregosa et al., 2011), Matplotlib (Hunter, 2007), MiniSom (Vettigli, 2013), Pandas (McKinney, 2010), and NumPy (Oliphant, 2015).

2.2.6.1. Feature selection

Table 2 lists all the features used in our model. A detailed description of each feature can be found in the ACE Level 2 documentation. To spread the data over a larger range of values in each component, we have used the logarithm of all the quantities, except of those marked with an asterisk in the table.

Features 11–15 contain an additional suffix, corresponding to a statistical operation performed on the corresponding feature. In our model we only include *range* operations, but we have provided our software with the ability to calculate also the mean, the standard deviation and the auto-correlation of quantities over a window of time of 6 h. This window allows to capture temporal (spatial) fluctuations in some of the solar wind parameters.

On the lower part of **Table 2** we present the range of dates used for the model. The same table also contains the hyper-parameters



selected to run the two models. The number of neurons per layer in the encoding half of the neural network is listed in the table.

2.2.6.2. Autoencoder architecture

We use a basic, fully connected feed-forward neural network for the encoding-decoding process. The neural network is symmetric in size but the weights of the encoder, W , and the decoder, W' , are not synchronized (see Equations 11, 12). Each layer is composed of a linear regressor, followed by a GELU activation function. The output layer of the network contains a linear regressor followed by a sigmoid activation function. The AE has been coded in python using the PyTorch framework (Paszke et al., 2019).

The final architecture of the AE and its hyper-parameters have been optimized automatically using the *Optuna* library (Akiba et al., 2019). We instructed this Hyper-Parameter Optimization (HPO) to select the optimal values for the following parameters, given the corresponding constraints:

- **Number of layers:** an integer between 2 and 6.
- **Number of neurons per layer:** it must be larger than 3 and smaller than the number of neurons in the previous encoder layer.
- **The neural network optimizer:** selected among Adam, Stochastic Gradient Descent, and RMSprop.
- **The learning rate:** a float value between 10^{-5} and 10^{-1} .

The automatic HPO is based on a technique called Tree-structured Parzen Estimator (TPE) (Bergstra et al., 2013), which uses Bayesian Optimization to minimize a target function, \mathcal{H} , provided by the user. We use the test loss of the AE as target function to be minimized.

The HPO performs a total of one thousand (1,000) different trials. However, to accelerate the optimization process, we built a smaller complementary data set. To avoid over-fitting on a subset of the original data we used the k -means algorithm to produce a representative sample of $m' = 4\sqrt{m}$ data points. This allows to explore a much broader set of hyperparameters in a short period of time. This artificial data set is then discarded and the AE is trained on the real data set.

The HPO selected the Adam optimizer (Kingma and Ba, 2014) for the gradient descent with a learning rate of 0.042. The total number of layers selected is 2, and the number of nodes in the bottleneck is 10. The loss function is the Mean Squared Error (MSE). We train the network for 500 epochs, after which no additional improvement in the loss function is observed. The full data set was randomly divided 50/50% between training and testing sets. We track the evolution of both data sets during training. We did not observe any variance or bias error.

The final architecture is trained using the full data set for 500 epochs. **Figure 5** shows the distribution of data in the original feature space, panel (A), and in the AE latent space, panel (B). The data in the original space contains extreme data points far from the mean value, and most features present a normal distribution. The combination of these two properties makes it difficult for any unsupervised clustering technique to separate points and accurately categorize different kinds of solar wind.

Panel (C) shows the error in the encoding-decoding procedure of the AE. It shows a histogram of the relative error, $E_r = \hat{X}/X -$

1, observed between the input data, X , and the decoded values, \hat{X} . A normal distribution function has been fitted to the values of the histogram. It shows that the relative error is centered near zero and its variance is around 1%.

2.2.6.3. Selection of parameters for the DSOM

In this manuscript we have introduced the use of the DSOMs for the classification of solar wind data. This technique requires the selection of four main Hyper-Parameters (HPs): the size of the lattice, $(L_x \times L_y)$, the constant learning rate, ϵ , and the elasticity, η . These last two parameters were chosen manually, while the lattice size was automatically selected by Hyper-Parameter Optimization (HPO) using *Optuna* (Akiba et al., 2019).

For the selection of the number of nodes in the lattice we propose the use of the objective function, \mathcal{H} , described in Equation (22):

$$\mathcal{H}(\sigma, \eta, L_x, L_y) = \frac{Q_E(\sigma, \eta, L_x, L_y)}{Q_0} + \alpha \frac{L_x}{m_{\max}} + \beta \frac{L_y}{n_{\max}} + \gamma \frac{L_x L_y}{\max(m_{\max}, n_{\max})} \quad (22)$$

where Q_E is the quantization error at the end of the training, Q_0 is a reference quantization error before training, L_x and L_y are the number of lattice nodes in each dimension, and m_{\max} and n_{\max} are the given maximum number of possible nodes. The weight factors α , β , and γ are used to impose restrictions on each term. We have fixed their value to $\alpha = \beta = 0.08$ and $\gamma = 0.6$. When a large number of nodes is available smaller values of Q_E are automatically obtained because the mean distance from the data set entries to the code words is reduced. The second and third terms on the RHS of \mathcal{H} leads the optimizer to reduce the number of nodes in the SOM. The squaring term $\gamma L_x L_y$ forces the map to be as squared as possible.

After a total of 500 trial runs of the model using different HPs, the optimizer selected the parameters presented in the lower section of **Table 2**. The optimization was accelerated using the same technique as in the optimization of the AE: we generated a reduced number of points using the k -means algorithm, with a total number of entries equal to one twentieth the size of the full data set, $m' = \frac{1}{20}m$.

The two remaining parameters of the DSOM, the elasticity $\eta = 3.0$ and the learning rate $\epsilon = 0.005$, have been manually selected. These two values control the speed at which the code words move toward the data entries, and the attraction between neighboring code words. It has been shown by Rougier and Boniface (2011) that high values of the elasticity, η , lead to tightly packed code words, while low values lead to loose connections. The elasticity takes in general values between 1 and 10. On the other hand, the learning rate indicates to the code words how fast they should move toward new incoming data. Very small learning rates could lead to very slow convergence to a solution, while very large values might produce code words that jump from value to value without converging to a global solution. The value of the learning rate can be set somewhere between 0.001 and 0.9.

Figure 6 shows how the elasticity and the learning rate can affect the convergence of the DSOM. In this figure we evaluate

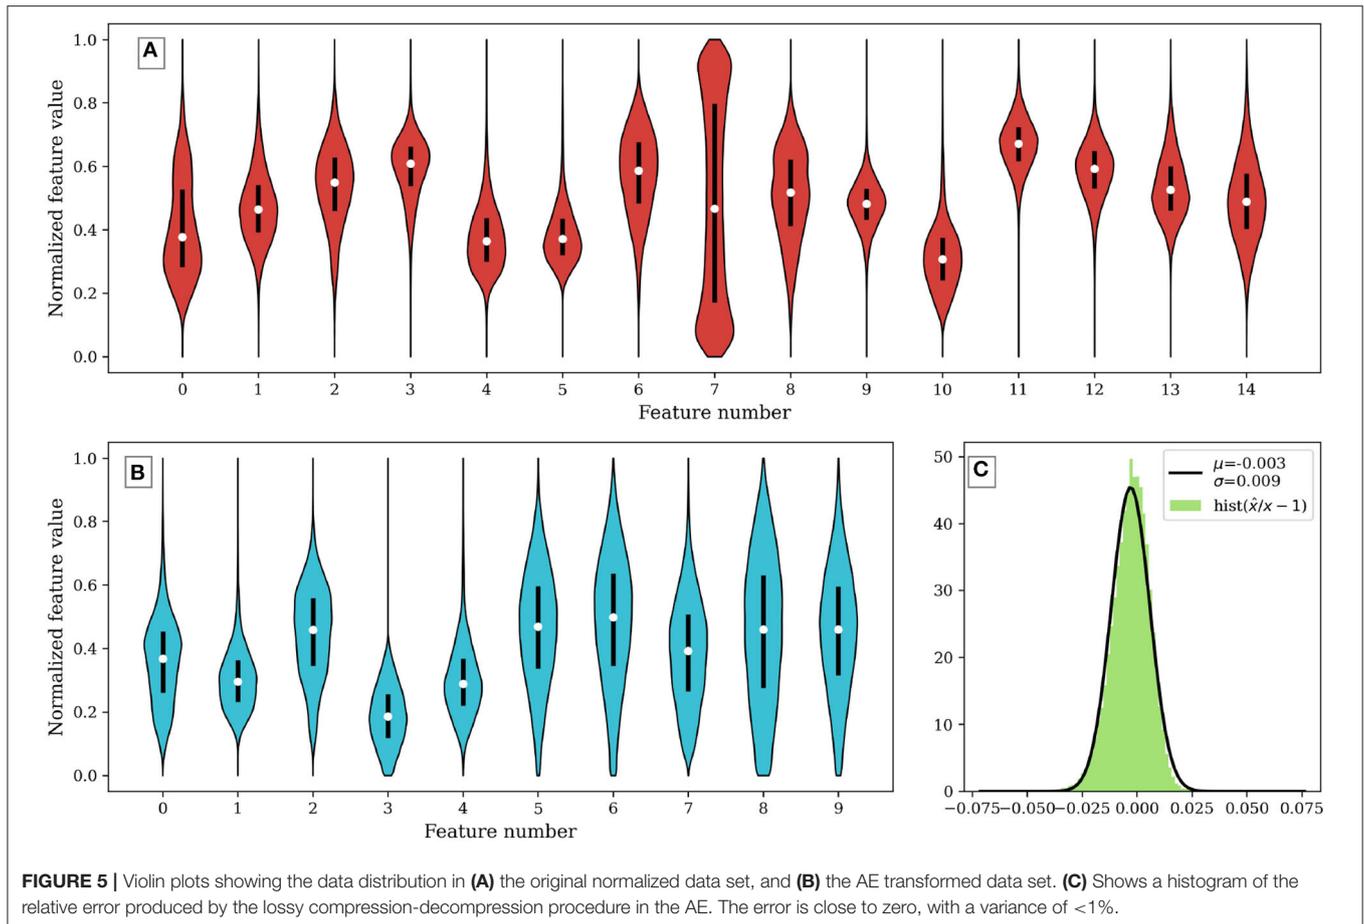


FIGURE 5 | Violin plots showing the data distribution in (A) the original normalized data set, and (B) the AE transformed data set. (C) Shows a histogram of the relative error produced by the lossy compression-decompression procedure in the AE. The error is close to zero, with a variance of <1%.

the effect of using different values of η and ϵ . Three different graphs are used to understand the evolution of the training and its convergence to a stable solution. The first row shows how the code words move away from their original position during the training; as the iterations advance the code words move until they find a stable location. It is clear that lower values of ϵ and η , as presented in the left panel of the first row, lead to very long convergence times. At the other extreme, very high values of the two parameters produce strong movements with a compact group of code words, leading to a non-converging solution.

In the second row of the same figure we show the distance traveled by the code words at each iteration of the training. In the best case scenario this distance is large at the beginning of the training and converges toward zero as the iterations pass. The third panel of this row shows how large values of η and ϵ produce solutions of the DSOM that do not converge.

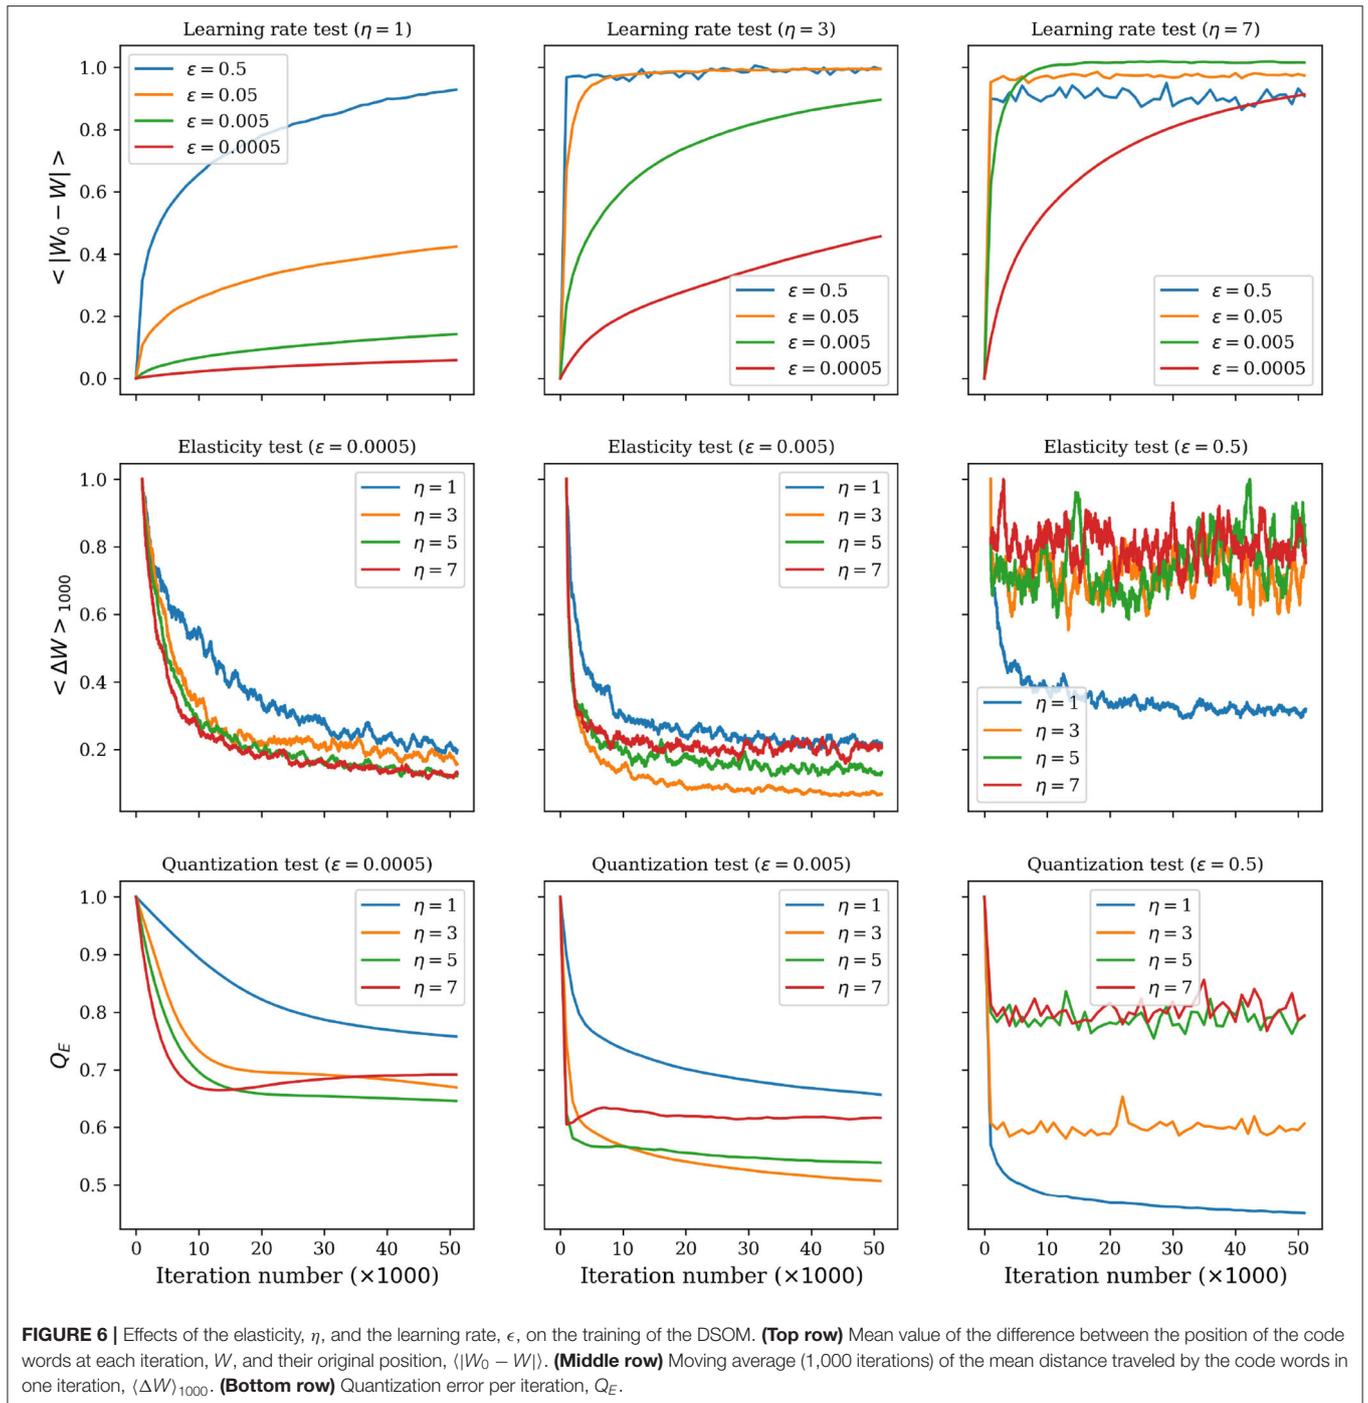
The third row of **Figure 6** shows the evolution of the quantization error (Equation 19). This value explains the compactness of the data points around the code words. Scattered points will show large Q_E , while dense clouds of points gathered around the code words will show low Q_E values. Once again in this last row we see that there is a compromise between a slow convergence with small values of ϵ , and large values of the two parameters that can lead to unstable solutions.

This figure also shows that, even if the DSOM is a dynamic technique that does not use a decay of the learning rate with time, it is a method that converges to a steady solution, if the parameters are properly selected.

2.2.6.4. Budget

Machine learning models require fine tuning of different parameters, from the selection and testing of multiple methods, to the parameterization of the final architecture. Dodge et al. (2019) suggests that every publication in machine learning should include a section on the budget used for the development and training of the method. The budget is the amount of resources used in the data processing, the selection of the model hyperparameters (HP), and its training.

The most time-consuming task in the present work has been the data preparation, the model setup and debugging and the writing of the SOM visualization routines. All the techniques described in the previous sections have been coded in python and are freely accessible in the repositories listed in the Data Availability Statement. We estimate the effort to bring this work from scratch to a total of two persons month. Of these, one person week was dedicated to the manual testing a selection of different model HPs (autoencoder architecture, feature selection, learning rates, initialization methods, number of epochs for training,



selection of data compression method, size of the time windows, etc.).

All classic clustering techniques presented in section 2.2.4 require only a few lines of code and can be trained in minutes on a mid-range workstation (e.g., Dell Precision T5600, featuring two Intel(R) Xeon(R) CPU E5-2643 0 @ 3.30GHz with four cores and eight threads each). The most time consuming tasks of our models are the training of the autoencoder (5% of the total run

time), the multiple passages of the clustering algorithms (15% of the run time), and the optimization of the hyper-parameters (80% of the run time). The training of the DSOM is performed in less than a minute.

For reference, the total run-time of our model is 30 min. The python scripts used do not contain any particular acceleration (e.g., using GPUs) or optimizations (e.g., using Numba), so there is large room for improvement of the computational efficiency.

3. RESULTS AND COMPARISONS

3.1. Interpretation of the DSOM Plots

When the DSOM method converges to a solution, each one of the code words is a representative of their N-dimensional neighborhood. We perform then a k-means clustering of the code words and apply the Kneedle method (Satopaa et al., 2011), presented in section 2.2.5.1, to select the final number of classes. Here, the automatic procedure selects a total of six classes, numbered from 0 to 5. The *Class Map* on the first panel of **Figure 7** shows that all nodes are organized in continuous groups.

The weights of the code words can be decoded and scaled to obtain the corresponding physical properties of the associated solar wind. These physical quantities are plotted in **Figure 7** for each one of the solar wind features.

Black continuous lines in the maps mark the boundary between different DSOM classes. All of the maps show uninterrupted smooth transitions between low and high values, without sudden jumps or incoherent color changes. Inside DSOM classes solar wind properties can present variations. This is an expected consequence of projecting 15 dimensions in a 2D lattice.

The most obvious class to identify is the DSOM class 0, with clear indications of coronal hole origin. It is characterized by very low values of the O^{7+}/O^{6+} and C^{6+}/O^{5+} ratios, associated with plasma originating from open magnetic field lines (Zhao et al., 2009; Stakhiv et al., 2016), high wind speed, low proton density, high absolute values of σ_c (a sign of Alfvénicity), high proton entropy, high proton temperature and moderately high values of Alfvén speed [associate by Xu and Borovsky (2015) with coronal holes].

The proton density has a very broad range of values for class 1. A close examination of the map of cross-helicity, σ_c , shows that this class also contains Alfvénic solar wind with both polarities. Class 1 also showcases high proton temperatures, high solar wind speeds, but average oxygen and carbon ionization ratio, and average iron charge. All these observations point toward solar wind originated at the boundary of coronal holes (Zhao et al., 2017).

Class 5 can be associated to transient event, such as ICME and ejecta. It presents the very high O^{7+}/O^{6+} ratio values that Zhao et al. (2009), Xu and Borovsky (2015), and Stakhiv et al. (2016) associate to CME plasma, and the low proton temperature values usually found in ICMEs. It is also characterized by the high (but, quite surprisingly, lower than for class 0) solar wind velocity, $\sigma_c \sim 0$ (Roberts et al., 2020), the high values of Alfvén speed which are usually associated to explosive transient activity (Xu and Borovsky, 2015).

Class 4 has similar properties as class 5 and can be mainly composed of transient events, but it also contains more Alfvénic plasma, and very high carbon charge state ratios, C^{6+}/C^{5+} . Fluctuations in this class are slightly less significant than the ones observed in class 5, except for jumps in the normal magnetic field, range B_n . These can point toward a class that contains magnetic clouds or Sector Boundary Crossing (SBC) events. Classes 4 and 5, identified as transients, remain rare, as clearly shown in **Figure 3C**.

At this point is important to remember that a different set of initial conditions or a different number of map nodes could lead to a slightly different repartition of the data, or to a different number of classes. However, points with similar properties will always remain topologically close and the interpretation of a different set of DSOM classes will lead to similar results. This is not necessarily the case with other unsupervised methods, like k-means, as the topological organization of the data is not maintained, so different runs can produce different results for which previous interpretations can not be re-cycled.

Class 2 and 3 are composed of slow, dense solar wind, the kind of wind that Zhao et al. (2017) associates to the *Quiet Sun* and that Xu and Borovsky (2015) associates to either Streamer Belt (SB) or Sector Reversal (SR) region plasma. As expected for the slow wind, the cross helicity is low, the proton temperature intermediate between the low values associated to ICMEs and the higher values observed in the fast wind, the proton entropy and the Alfvén speed are low (Xu and Borovsky, 2015). The high O^{7+}/O^{6+} and C^{6+}/C^{5+} ratios (lower only to the values associates to class mappable to transient events, class 4 and 5), point to plasma originating in closed field lines (Zhao et al., 2009; Stakhiv et al., 2016). Of the two classes, class 2 is characterized by lower wind speed, higher density, lower proton temperature, lower entropy.

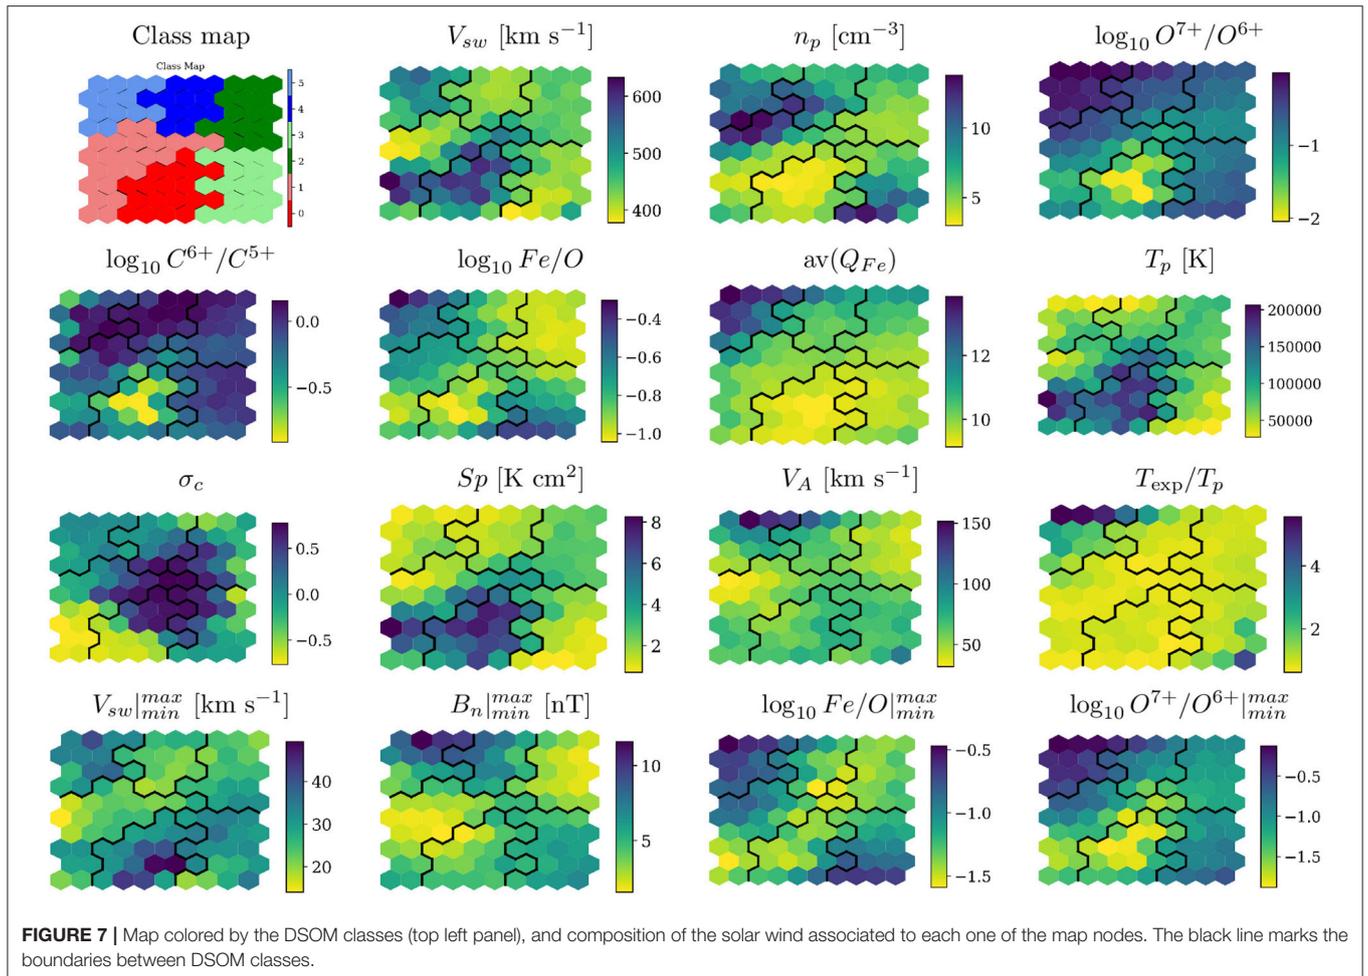
In summary we can group our classes on three major categories: CH wind (classes 0 and 1, colored in red), quiet or transitional wind (classes 2 and 3, colored in green), and transients (classes 4 and 5, colored in blue).

3.2. Verification of the DSOM Classes

In addition to the interpretation of the maps presented in the previous section, we have extracted histograms of the occurrence frequency of O^{7+}/O^{6+} ratio (**Figure 8**) and proton speed, V_{sw} (**Figure 9**). The panels in the figures contain the histograms for six (6) different categorizations: k-means (AE), k-means (KPCA), BGM (AE), BGM (KPCA), DSOM, and the X15 classification. All the histograms have been normalized row by row (class by class), following the work done by Zhao et al. (2017). This representation of the data is inspired by Figure 5 of that paper, where the authors showed an important overlapping among different solar wind classes, and a bi-modal velocity distribution for coronal hole wind including an important population of slow wind.

The assignment of class numbers by the clustering algorithms is random. We have sorted the classes so they present an ascending value of the O^{7+}/O^{6+} ratio in **Figure 8**. It has been shown that solar wind originated in *Coronal Holes* present very low values of the O^{7+}/O^{6+} , while at the other extreme transient events present very high O^{7+}/O^{6+} ratios (Zhao et al., 2009, 2017; Stakhiv et al., 2016). **Figure 8** confirms the class identification we presented in the previous section.

von Steiger and Zurbuchen (2015) and Bloch et al. (2020) examine the O^{7+}/O^{6+} ratio in Ulysses data, which include abundant measures of wind originating from the polar CHs. Our data is composed of ACE observations from the ecliptic plane. For this reason, in all different classifications in **Figure 8**,



including the X15 empirical categorization, class 0 does not reach $\log_{10} O^{7+}/O^{6+} \approx -2$, where the peak of points is observed in publications using Ulysses data.

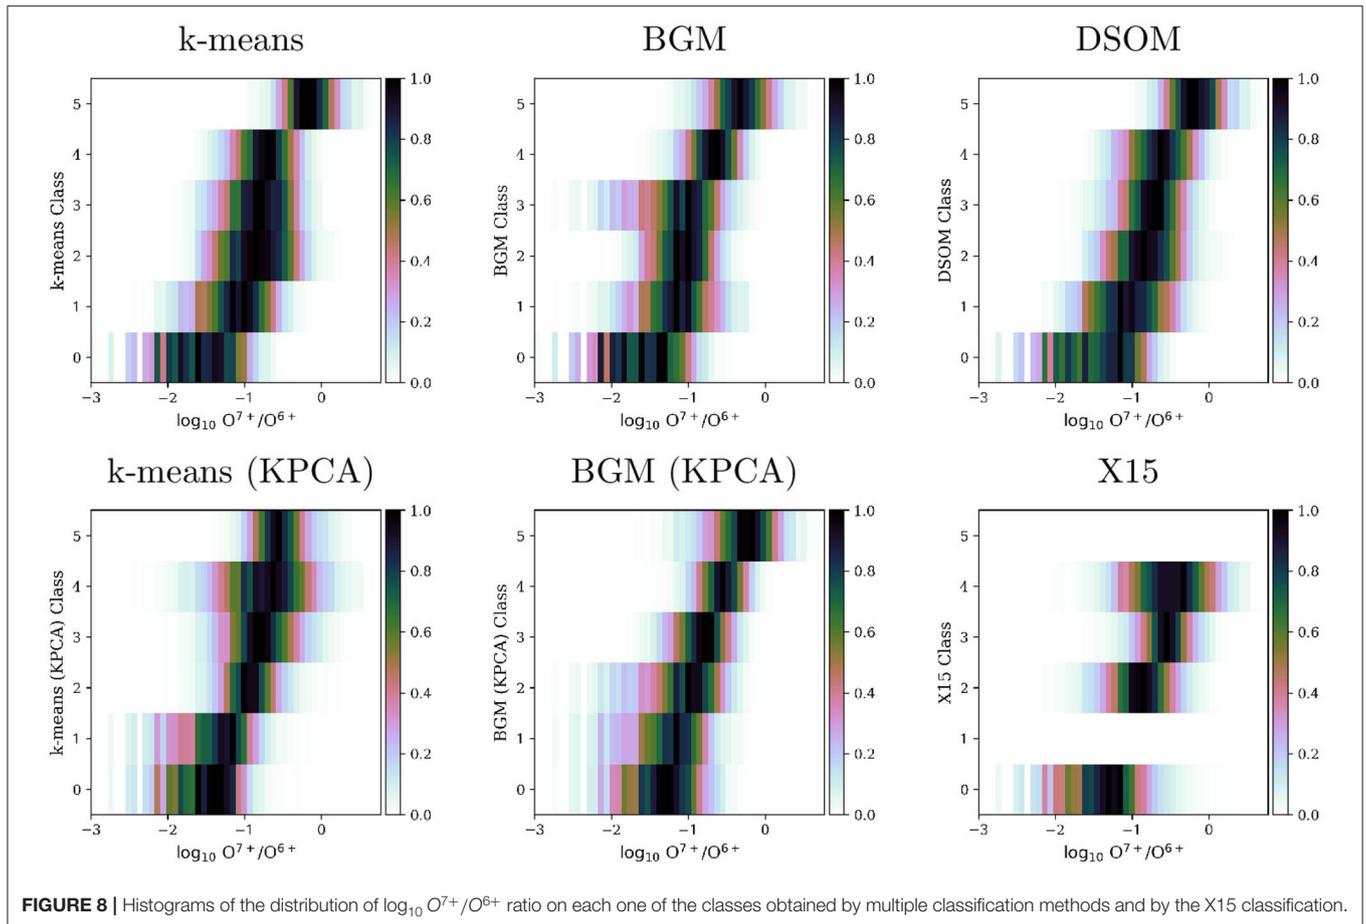
In our data set, the majority of points can be mapped to *Quiet Sun* (QS), conditions, i.e., slow solar wind. Even in these conditions, the DSOM method is able to sample and distribute enough data to each one of the classes. The BGM method applied to Kernel PCA transformed data also provide a good sample of the different classes, in particular for transient solar wind (classes 4 and 5). The X15 classification was designed with clear boundaries in O^{7+}/O^{6+} , for this reason the differences among the four classes is clear in the histograms. However, this observation contradicts the foot point back tracing performed by Zhao et al. (2017): X15 shows almost no overlap in the distribution functions between the different classes, while the back tracing shows important overlaps. We express caution in the use of this classification to train any type of supervised machine learning technique, or in the evaluation of the accuracy of unsupervised techniques.

Figure 9 shows how velocity is distributed among the different classes for each unsupervised classification method, and for the X15 categories. Zhao et al. (2017) remarks that the different classes are more difficult to identify using the solar wind speed

histograms. We verify in these plots that three conditions are satisfied: (1) the classes we associate to the QS (class 2 and 3 in the DSOM classification) are associated to low velocity regions (Neugebauer et al., 2002), (2) high oxygen state ratios are associated with low solar wind speeds, and (3) CH wind has a highly spread velocity distribution, with two possible peaks around 400 and 600 km/s (Zhao et al., 2017).

The fact that class 0 and 1, that we associate to wind of CH origin, contains slow wind data points is particularly significant. D'Amicis and Bruno (2015) has provided proof of the presence, at 1 AU, of highly Alfvénic slow wind originating from the boundaries of coronal holes. This slow, Alfvénic wind has the same composition signature and high cross helicity that characterized the *classic* fast Alfvénic wind of CH origin, but presents lower speed and lower proton temperature. This way of visualizing our results seems to suggest that slow Alfvénic wind is classified together with fast Alfvénic wind in the classes that we associate to CH origin.

Figure 9 shows that CH wind in the *k*-means and DSOM plots present a broad range of speeds, with a bimodal distribution. The BGM (KPCA) method separates these two populations in two different classes (0 and 1). the *k*-means (KPCA) method



differentiates the fast solar wind, in the first two classes, from the slow wind in the remaining classes.

Balancing the results from **Figures 8, 9** we conclude that the BGM (KPCA) and the DSOM are the techniques that approach the most the direct observations of the solar wind origins obtained by Zhao et al. (2017). The X15 model creates a very sharp separation of solar wind types, with fast winds clearly segregated in class 0, slow winds in classes 1 and 2, and transients in class 3. The X15 model does not recognize that plasma of CH origin also contains an important population of slow winds.

3.3. Hit Maps of Empirical Classifications

Another advantage of the SOM/DSOM method is that it can be used to visualize additional hidden statistics. **Figure 10** shows what nodes are activated by the Z09 and X15 classes. To perform this analysis, instead of using the full data set, we extract three subsets corresponding to the entries categorized as CH, ICME, and NCH wind in the Z09, and CH, SB, SR, and ICME in the X15 catalogs. Each one of these three (four) subsets is passed through the DSOM model and we observe how each one activates the nodes.

We see that CH wind, in column 1 of the figure, activates very similar nodes for both classifications, in classes 0 and 1. Most of the hits are located on nodes where the absolute value of the

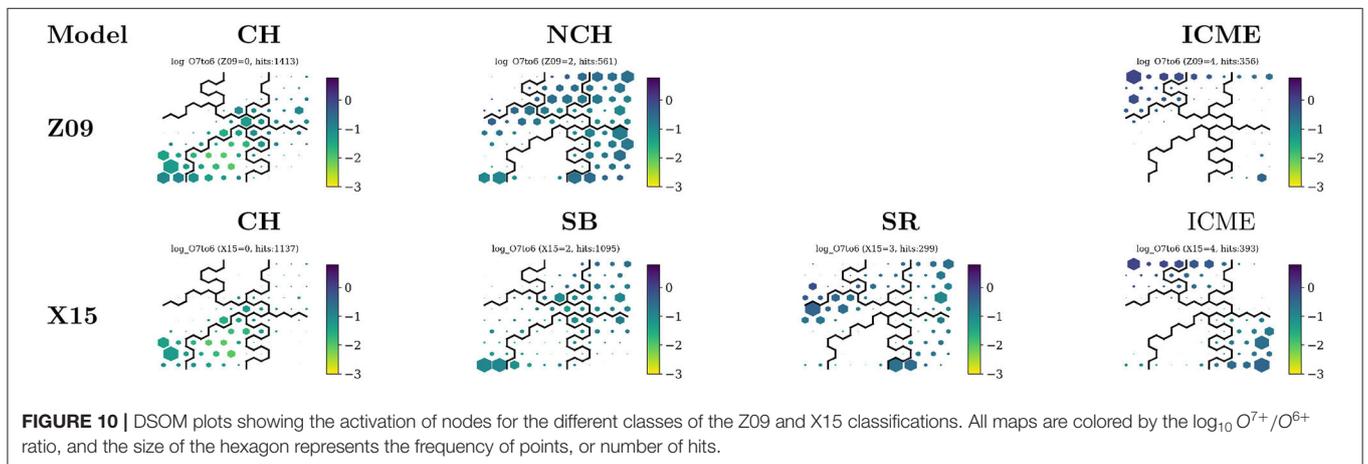
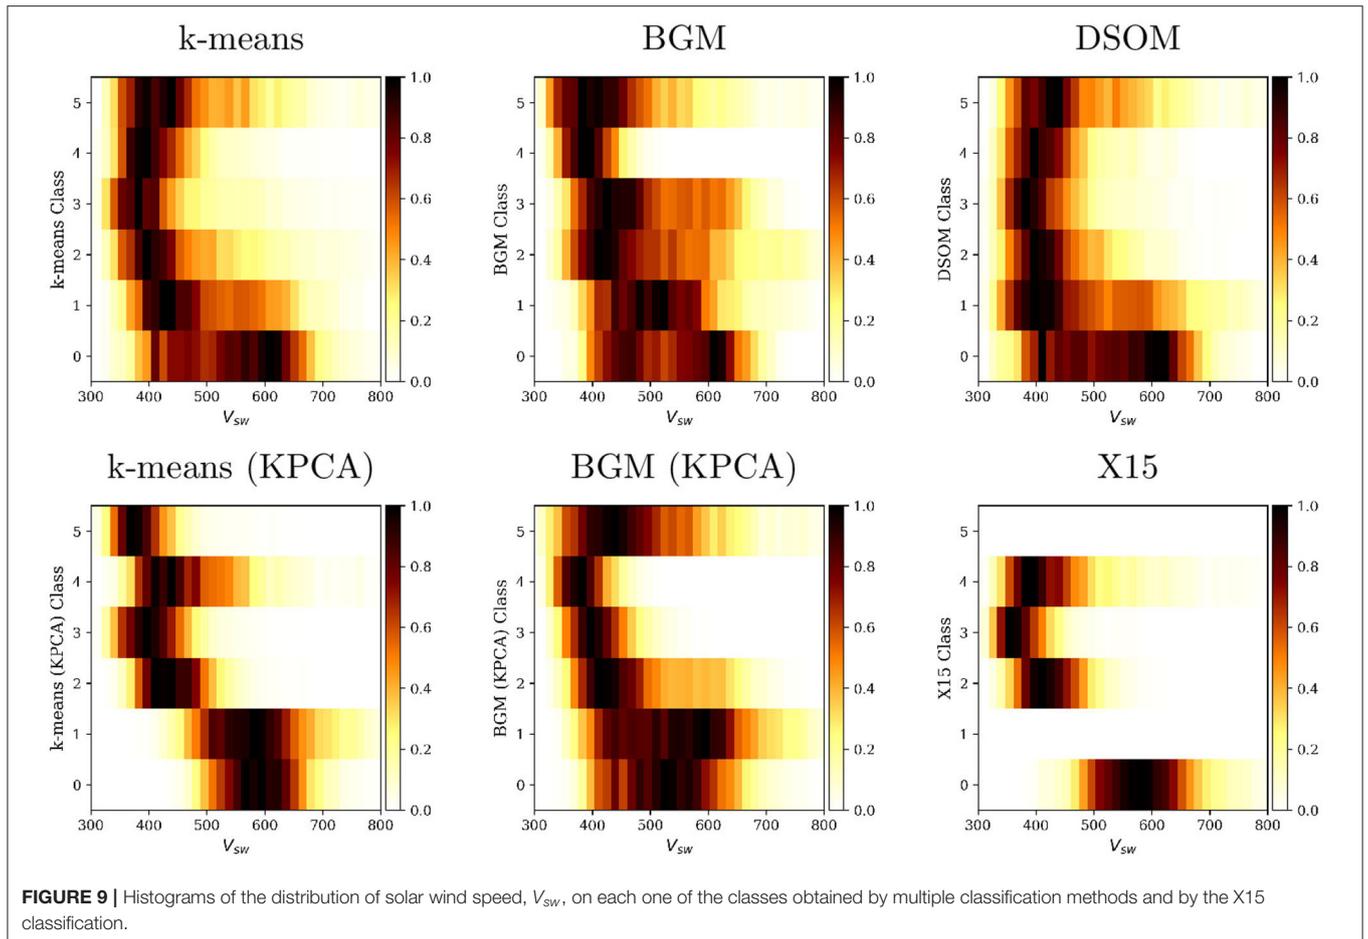
cross-helicity σ_c is the largest, i.e., in regions of open field lines associated with coronal holes.

NCH wind from the Z09 classification is distributed over classes 2, 3, and 4, but also includes a node from class 1 characterized by an extremely negative cross-helicity. The same zone is activated by the SB class from X15. The two affected nodes also feature a very low T_{exp}/T_p . The X15 model splits solar wind points using hyperplanes in a three-dimensional space composed by S_p , O^{7+}/O^{6+} , and T_{exp}/T_p . None of those planes cuts the points in the T_{exp}/T_p dimension (Xu and Borovsky, 2015). However, in our maps this dimension seems to play an important role in the separation between quiet and CH winds.

The X15 Sector Reversal (SR) class activates nodes at the boundaries of classes 1, 2, and 3. These nodes separate the quiet sun from the coronal hole wind, and coronal holes to transients. It also contains a large population of slow quiet solar wind.

Finally transients, in both the Z09 and X15 categorizations, are associated to our class 4 and 5. However, a large portion of the X15 transients is associated to class 3 of the DSOM, particularly in nodes showing low proton temperatures and specific entropy S_p , characteristics of ICMEs.

Figure 10 shows also that, on average, the values of the O^{7+}/O^{6+} ratio do not change radically among the nodes, except for small variations in the CH and the ICME classes of X15.



3.4. Quantitative Comparison With Empirical Classifications

We have included a *Matching Matrix* in **Table 3** showing the frequency of occurrences of our model with respect to the Z09 and X15 classifications. Bold numbers in the table mark the highest common frequency and regular fonts mark the second highest frequency for

each one of the columns. Matching matrices must not be confused with confusion matrices, as the later imply that there is a *ground truth*. Matching matrices are used in unsupervised learning to compare the frequency of occurrence of classes between models, so we can not perform additional metrics, like accuracy, precision, sensitivity, or specificity.

TABLE 3 | Matching matrix comparing the DSOM, X15, and Z09 classifications.

		DSOM Class						Z09		
		0	1	2	3	4	5	CH	NCH	TR
X15	CH	7,727	3,994	3,330	125	<u>502</u>	<u>47</u>	14,273	993	459
	SB	<u>7,423</u>	6,295	7,194	<u>916</u>	950	138	<u>12,904</u>	9,244	<u>768</u>
	SR	3,233	392	1,404	584	434	<u>157</u>	877	<u>4,747</u>	580
	TR	1,263	575	387	2,373	403	1,362	1,343	3,011	2,009
DSOM Class	0							13,357	<u>5,966</u>	323
	1							9,040	<u>2,213</u>	3
	2							6,173	<u>5,789</u>	353
	3							637	2,848	513
	4							170	<u>987</u>	1,132
	5							20	<u>192</u>	1,492

Values in bold (underline) font represent the highest (second highest) frequency for each column (row) in the top (bottom) half of the table.

In this matrix we see that CH and SB categories from the X15 classification are mostly associated with classes 0, 1, 2, and 4 in the DSOM model, while TR winds are associated with classes 3 and 5. No particular class is clearly associated with SR winds, but the highest frequency is observed for class 0.

CH in the Z09 classification are accurately associated with classes 0 and 1, but a big part of the NCH wind is also grouped in class 0. Transients are correctly distributed among classes 4 and 5 of the DSOM.

We highlight that the X15 and Z09 models, the two classifications most used for the verification of machine learning results (see Camporeale et al., 2017; Li et al., 2020), are not fully compatible among themselves. A large number of CH winds from the Z09 classification is associated with SB winds in the X15 classification, and a considerable number of transients are cataloged as sector boundary crossings (SB).

3.5. Time Series Comparisons

A complementary method to compare the different classification techniques is to visually inspect windows of time and check, with the help of a human expert, that the time series are in agreement with the previous analysis. **Figure 11** shows, in two columns, two windows of time of 4 months. The left column contains a high solar activity period, from May 2003 to September 2003, and the right column contains a period of low solar activity, between January 2008 and May 2008. Each one of the eight (8) rows contains a plot of the solar wind speed colored by a different classification method, from empirical models (Z09, vS15, and X15) to unsupervised methods (*k*-means, *k*-means, BGM, and DSOM). The colors of the empirical methods in the time series correspond to the labels assigned in **Table 1**, and the colors of the models were all assigned by manually ordering the classes following the frequency $\log_{10}(O^{7+}/O^{6+})$, from low values (low category number) to high values (high category number).

In the same figure vertical gray zones correspond to Richardson and Cane ICME catalog entries (Richardson and Cane, 2012), and vertical lines to entries in the UNH and CfA catalogs.

It is clear that among the empirical models, the vS15, based on observations by the Ulysses mission, is the most restrictive in the selection of CH origin winds, however during the plotted quiet time in the right column, which corresponds to the declining phase of the solar cycle, a significant part of the solar wind originates in coronal holes, and in fact High Speed Stream and Corotating Interaction regions, associated to wind of CH origin, are the main driver of geomagnetic activity during the declining phase of the cycle (Tsurutani et al., 2006; Innocenti et al., 2011). During both solar activity windows the Z09 and X15 models assign an important number of observations to coronal holes. von Steiger and Zurbuchen (2015) shows that the threshold used in the Z09 classification to identify coronal holes is not accurate and can misclassify NCH as CH. Both models accurately identify transients in the data. Quiet solar wind is more clearly visible during the low solar activity window in the X15 model.

k-means (KPCA) and BGM (KPCA) correctly classify CH origin winds (classes 0 and 1). A clear transition between class 1, CH wind, and class 3 can be observed on both panels. Transients are also well-captured with classes 4 and 5. On the other hand, classifications based on the *k*-means (AE) and BGM (AE), do not show high accuracy in these two windows of time, but are able to detect transients. These two methods show difficulties in discerning QS winds from CH solar winds.

The DSOM model shows good performances. The two classes associated with CH origin wind, classes 0 and 1, are more restrictive than the Z09 and X15 classes. Classes 4 and 5 distinguish between two different types of transients. ICMEs in these time windows are mainly associated with class 5, except for transients observed around 2003-05-20 and 2003-06-15. The model also detects a very slow transient around 2003-07-10. The 27 days solar period is also evident on the oscillations of the solar wind speed and the periodic nature of the solar wind types. In the low solar activity window the solar wind is more homogeneous and shows mainly CH and QS origin winds, as expected (Tsurutani et al., 2006).

Different classification methods lead to different classes with different properties. Roberts et al. (2020) performed detailed descriptions of the categorized solar wind classes based on the mean values observed in each subset of points. Zhao et al. (2017) shows that it is important to look at the frequency distribution and not only the mean. Our model shows that some features can present very large distributions inside a single class, even multiple peaks, as is the case of the solar wind speed for the CH classes.

We will perform further refinements of the model and its interpretation in a future work. These preliminary results show the great potential of the techniques introduced in this paper. DSOMs show the variability of solar wind and how it can be visually characterized. The DSOM is a helpful guide in the study of the different types of solar wind, but is not necessarily an objective, unbiased and final classification method. In our current understanding, the main factor that determines classification results is the choice of the solar wind parameters used in the DSOM training. Choosing parameters that, according to previous studies and our physical understanding of the wind, can discriminate between specific wind types can guide the classification results. On the other hand, the possibility exists that

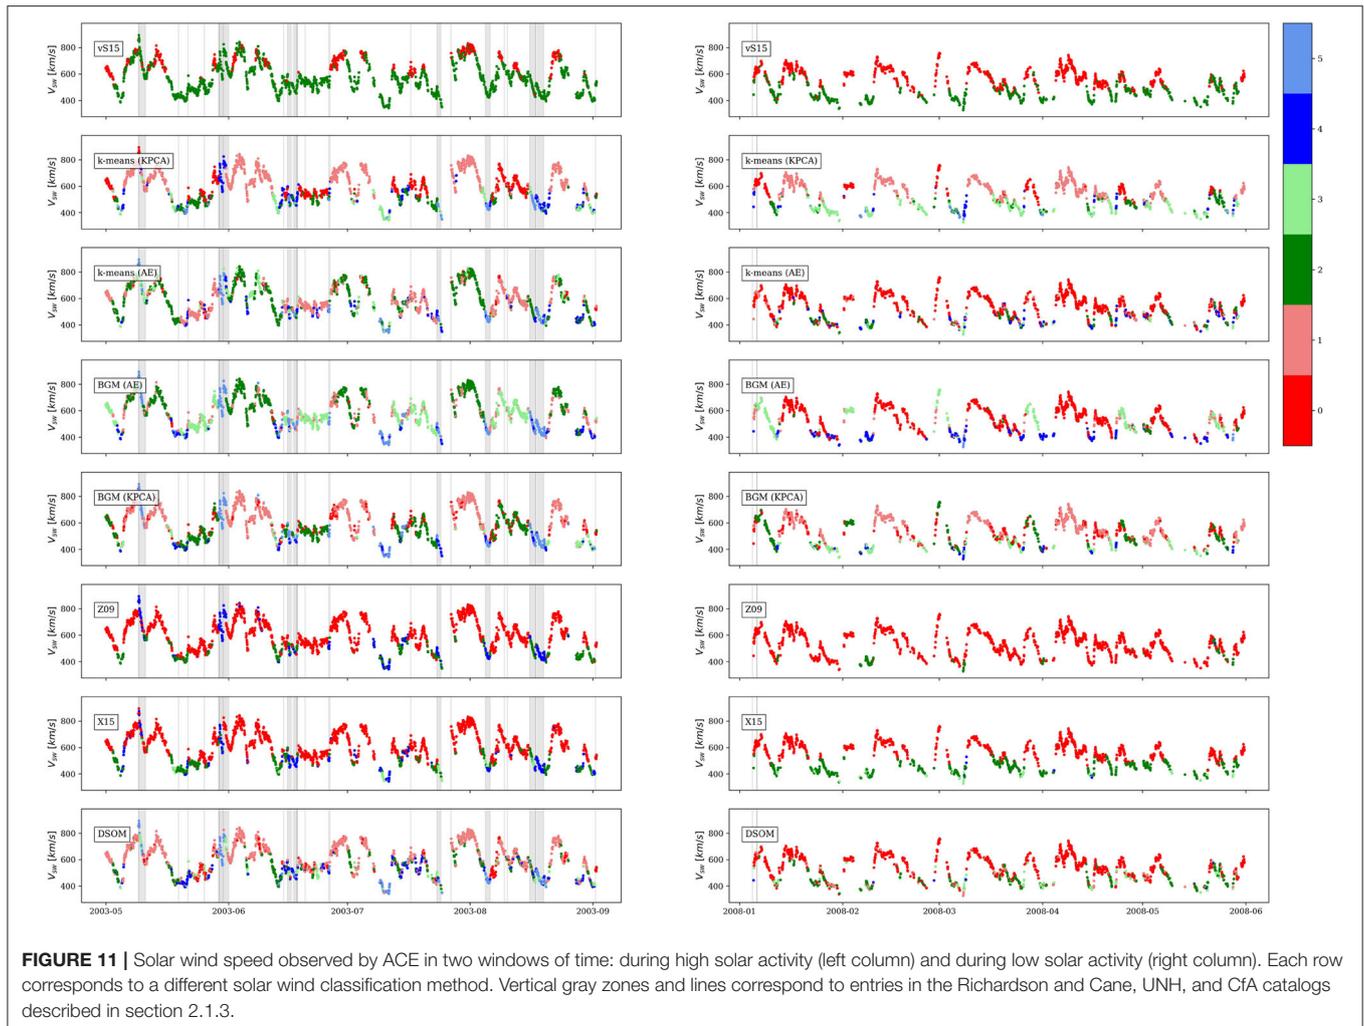


FIGURE 11 | Solar wind speed observed by ACE in two windows of time: during high solar activity (left column) and during low solar activity (right column). Each row corresponds to a different solar wind classification method. Vertical gray zones and lines correspond to entries in the Richardson and Cane, UNH, and CfA catalogs described in section 2.1.3.

an unsupervised classification methods, such as the one used here will highlight the presence of solar wind types that could warrant future physical investigation. DSOMs open the possibility for a fast visual characterization of large and complex data sets.

4. DISCUSSION

In this paper we show how the categorization of solar wind can be informed by classic unsupervised clustering methods and Self-Organizing Maps (SOM). We demonstrate that a single technique used in isolation can be misleading in the interpretation of automatic classifications. We show that it is important to examine the SOM lattices, in conjunction with solar wind composition and velocity distributions, and time series plots. Thanks to these tools we can differentiate classes associated with known heliospheric events.

We are convinced that basic unsupervised clustering techniques will have difficulties in finding characteristic solar wind classes when they are applied to unprocessed data. A

combination of feature engineering, non-linear transformations and SOM training leads to a more appropriate segmentation of the data points.

The classification of the solar wind also depends on the objectives that want to be attained: if the goal is to classify the solar wind to study its origin on the Sun, features related to solar activity must be included in the model; however, if the goal is to identify geoeffectiveness, other parameters should be added to the list of features, including geomagnetic indices.

In this work we have presented a first test of the capabilities of the SOMs for the analysis of data from a full solar cycle. Due to the extent of the work done, in this paper we introduce all the methods and techniques developed, but we leave for a future publication a more refined selection of all the model parameters, and the corresponding interpretation of the solar wind classification.

Finally, we advocate for the creation of a catalog of foot point locations for every solar mission, that connect solar wind observations to points on the solar surface. Due to the uncertainty of on the exact foot point, such catalog should be composed of

a set of probabilities for each possible solar origin. This ground truth will vastly improve the efficacy of our classification models, which in turn can be used to reduce the initial uncertainties of the catalog.

All the tools and the techniques presented here can be applied to any other data set consisting of large amounts of points with a fixed number of properties. All the software and the data used in this work are freely available for reproduction and improvement of the results presented above.

DATA AVAILABILITY STATEMENT

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://github.com/murci3lag0/swinsom>.

AUTHOR CONTRIBUTIONS

JA created the software used in this work, built the data sets, and wrote the manuscript. RD provided the important insights into the use of the machine learning techniques, and performed

revisions of the different drafts. MEI gathered the information from external collaborators, provided insights into the data usage, and proofread the different drafts. GL supervised the work. All authors contributed to the manuscript revision, read, and approved the submitted version.

FUNDING

The work presented in this manuscript has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 754304 (DEEP-EST, www.deep-projects.eu), and from the European Union's Horizon 2020 research and innovation programme under grant agreement No 776262 (AIDA, www.aida-space.eu).

ACKNOWLEDGMENTS

The authors would like to acknowledge the helpful advice and suggestions by Olga Panasenco (UCLA), Raffaella D'Amicis (INAF-IAPS), and Aaron Roberts (NASA-GSFC). We thank the teams of the ACE SWEPAM/SWICS/MAG/EPAM instruments and the ACE Science Center for providing the ACE data.

REFERENCES

- Adhikari, L., Zank, G. P., Zhao, L.-L., Kasper, J. C., Korreck, K. E., Stevens, M., et al. (2020). Turbulence transport modeling and first orbit parker solar probe (PSP) observations. *Astrophys. J. Suppl. Ser.* 246:38. doi: 10.3847/1538-4365/ab5852
- Akiba, T., Sano, S., Yanase, T., Ohta, T., and Koyama, M. (2019). "Optuna: a next-generation hyperparameter optimization framework," in *Proceedings of the ACM SIGKDD International Conference on Knowledge Discovery and Data Mining* (Anchorage, AK), 2623–2631. doi: 10.1145/3292500.3330701
- Bale, S. D., Badman, S. T., Bonnell, J. W., Bowen, T. A., Burgess, D., Case, A. W., et al. (2019). Highly structured slow solar wind emerging from an equatorial coronal hole. *Nature* 576, 237–242. doi: 10.1038/s41586-019-1818-7
- Bandyopadhyay, R., Goldstein, M. L., Maruca, B. A., Matthaeus, W. H., Parashar, T. N., Ruffolo, D., et al. (2020). Enhanced energy transfer rate in solar wind turbulence observed near the sun from Parker solar probe. *Astrophys. J. Suppl. Ser.* 246:48. doi: 10.3847/1538-4365/ab5dae
- Bergstra, J., Yamins, D., and Cox, D. (2013). "Making a science of model search: hyperparameter optimization in hundreds of dimensions for vision architectures," in *Proceedings of the 30th International Conference on Machine Learning Volume 28 of Proceedings of Machine Learning Research*, eds S. Dasgupta and D. McAllester (Atlanta, GA: PMLR), 115–123.
- Bishop, C. M. (2006). *Machine Learning and Pattern Recognition*. Information Science and Statistics. Heidelberg: Springer.
- Bloch, T., Watt, C., Owens, M., McInnes, L., and Macneil, A. R. (2020). Data-driven classification of coronal hole and streamer belt solar wind. *Sol. Phys.* 295:41. doi: 10.1007/s11207-020-01609-z
- Burlaga, L., Sittler, E., Mariani, F., and Schwenn, R. (1981). Magnetic loop behind an interplanetary shock: Voyager, Helios, and IMP 8 observations. *J. Geophys. Res. Space Phys.* 86, 6673–6684. doi: 10.1029/JA086iA08p06673
- Camporeale, E., Carè, A., and Borovsky, J. E. (2017). Classification of solar wind with machine learning. *J. Geophys. Res. Space Phys.* 122, 10910–10920. doi: 10.1002/2017JA024383
- Cane, H. V., and Richardson, I. G. (2003). Interplanetary coronal mass ejections in the near-Earth solar wind during 1996–2002. *J. Geophys. Res. Space Phys.* doi: 10.1029/2002JA009817. [Epub ahead of print].
- D'Amicis, R., and Bruno, R. (2015). On the origin of Highly Alfvénic slow solar wind. *Astrophys. J.* 805:84. doi: 10.1088/0004-637X/805/1/84
- Dodge, J., Gururangan, S., Card, D., Schwartz, R., and Smith, N. A. (2019). Show your work: improved reporting of experimental results. *arXiv [Preprint]* arXiv:1909.03004.
- Dupuis, R., Goldman, M. V., Newman, D. L., Amaya, J., and Lapenta, G. (2020). Characterizing magnetic reconnection regions using gaussian mixture models on particle velocity distributions. *Astrophys. J.* 889:22. doi: 10.3847/1538-4357/ab5524
- Eselevich, V. G., and Filippov, M. A. (1988). An investigation of the heliospheric current sheet (HCS) structure. *Planet. Space Sci.* 36, 105–115. doi: 10.1016/0032-0633(88)90046-3
- Ester, M., Kriegel, H.-P., Sander, J., and Xu, X. (1996). "A density-based algorithm for discovering clusters in large spatial databases with noise," in *KDD'96: Proceedings of the Second International Conference on Knowledge Discovery and Data Mining* (Portland, OR), Vol. 96, 226–231.
- Feldman, U., Landi, E., and Schwadron, N. A. (2005). On the sources of fast and slow solar wind. *J. Geophys. Res. Space Phys.* 110. doi: 10.1029/2004JA010918
- Fisk, L., and Lee, M. (1980). Shock acceleration of energetic particles in corotating interaction regions in the solar wind. *Astrophys. J.* 237, 620–626. doi: 10.1086/157907
- Fu, H., Li, B., Li, X., Huang, Z., Mou, C., Jiao, F., et al. (2015). Coronal sources and *in situ* properties of the solar winds sampled by ACE during 1999–2008. *Sol. Phys.* 290, 1399–1415. doi: 10.1007/s11207-015-0689-9
- Garrard, T. L., Davis, A. J., Hammond, J. S., and Sears, S. R. (1998). *The ACE Science Center BT—The Advanced Composition Explorer Mission*. Dordrecht: Springer Netherlands.
- Gloeckler, G., Cain, J., Ipavich, F. M., Tums, E. O., Bedini, P., Fisk, L. A., et al. (1998). *Investigation of the Composition of Solar and Interstellar Matter Using Solar Wind and Pickup Ion Measurements With SWICS and SWIMS on the Ace Spacecraft BT—The Advanced Composition Explorer Mission*. Dordrecht: Springer Netherlands.
- Gold, R. E., Krimigis, S. M., Hawkins, S. E., Haggerty, D. K., Lohr, D. A., Fiore, E., et al. (1998). *Electron, Proton, and Alpha Monitor on the Advanced*

- Composition Explorer Spacecraft BT-The Advanced Composition Explorer Mission. Dordrecht: Springer Netherlands.
- Goodfellow, I. J., Pouget-Abadie, J., Mirza, M., Xu, B., Warde-Farley, D., Ozair, S., et al. (2014). Generative Adversarial Networks *arXiv [Preprint]* arXiv:1406.2661.
- Gopalswamy, N., Hanaoka, Y., Kosugi, T., Lepping, R. P., Steinberg, J. T., Plunkett, S., et al. (1998). On the relationship between coronal mass ejections and magnetic clouds. *Geophys. Res. Lett.* 25, 2485–2488. doi: 10.1029/98GL50757
- Gosling, J. T., Teh, W. L., and Eriksson, S. (2010). A torsional Alfvén wave embedded within a small magnetic flux rope in the solar wind. *Astrophys. J. Lett.* 719, 36–40. doi: 10.1088/2041-8205/719/1/L36
- Habbal, S. R., Woo, R., Fineschi, S., O’Neal, R., Kohl, J., Noci, G., et al. (1997). Origins of the slow and the ubiquitous fast solar wind. *Astrophys. J.* 489, L103–L106. doi: 10.1086/310970
- Heidrich-Meisner, V., and Wimmer-Schweingruber, R. F. (2018). “Chapter 16—solar wind classification via k-means clustering algorithm,” in *Machine Learning Techniques for Space Weather*, eds E. Camporeale, S. Wing, and J. R. Johnson (Elsevier), 397–424.
- Hinton, G. E., and Salakhutdinov, R. R. (2006). Reducing the dimensionality of data with neural networks. *Science* 313, 504–507. doi: 10.1126/science.1127647
- Hunter, J. D. (2007). Matplotlib: a 2D graphics environment. *Comput. Sci. Eng.* 9, 90–95. doi: 10.1109/MCSE.2007.55
- Innocenti, M. E., Lapenta, G., Vršnak, B., Crespon, F., Skandrani, C., Temmer, M., et al. (2011). Improved forecasts of solar wind parameters using the Kalman filter. *Space Weather* 9. doi: 10.1029/2011SW000659
- Kingma, D. P., and Ba, J. (2014). Adam: A Method for Stochastic Optimization *arXiv [Preprint]* arXiv:1412.6980.
- Kingma, D. P., and Welling, M. (2013). Auto-Encoding Variational Bayes *arXiv [Preprint]* arXiv:1312.6114.
- Li, H., Wang, C., Tu, C., and Xu, F. (2020). Machine learning approach for solar wind categorization. *Earth Space Sci.* 7:e2019EA000997. doi: 10.1029/2019EA000997
- Liang, J., and Liu, R. (2015). “Stacked denoising autoencoder and dropout together to prevent overfitting in deep neural network,” in *2015 8th International Congress on Image and Signal Processing (ICISP)* (Shenyang), 697–701. doi: 10.1109/CISP.2015.7407967
- Lloyd, S. (1982). Least squares quantization in PCM. *IEEE Trans. Inform. Theory* 28, 129–137. doi: 10.1109/TIT.1982.1056489
- Magyar, N., Van Doorslaere, T., and Goossens, M. (2019). The nature of Elsässer variables in compressible MHD. *Astrophys. J.* 873:56. doi: 10.3847/1538-4357/ab04a7
- Matteini, L., Hellinger, P., Landi, S., Trávníček, P. M., and Velli, M. (2012). Ion kinetics in solar wind: coupling global expansion to local microphysics. *Space Sci. Rev.* 172, 373–396. doi: 10.1007/s11214-011-9774-z
- McComas, D. J., Bame, S. J., Barker, P., Feldman, W. C., Phillips, J. L., Riley, P., et al. (1998). *Solar Wind Electron Proton Alpha Monitor (SWEPAM) for the Advanced Composition Explorer BT-The Advanced Composition Explorer Mission*. Dordrecht: Springer Netherlands.
- McKinney, W. (2010). “Data structures for statistical computing in Python,” in *Proceedings of the 9th Python in Science Conference*, eds S. van der Walt and J. Millman (Austin, TX), 56–61. doi: 10.25080/Majora-92bf1922-00a
- Neugebauer, M., Lieber, P. C., Smith, E. J., Skoug, R. M., and Zurbuchen, T. H. (2002). Sources of the solar wind at solar activity maximum. *J. Geophys. Res. Space Phys.* 107, SSH 13-1–SSH 13-15. doi: 10.1029/2001JA000306
- Neugebauer, M., and Snyder, C. W. (1966). Mariner 2 observations of the solar wind: I. Average properties. *J. Geophys. Res.* 71, 4469–4484. doi: 10.1029/JZ071i019p04469
- Oliphant, T. E. (2015). *Guide to NumPy, 2nd Edn*. North Charleston, SC: CreateSpace Independent Publishing Platform.
- Owens, M. J. (2016). Do the legs of magnetic clouds contain twisted flux-rope magnetic fields? *Astrophys. J.* 818:197. doi: 10.3847/0004-637X/818/2/197
- Parker, E. N. (1958). Interaction of the solar wind with the geomagnetic field. *Phys. Fluids* 1, 171–187. doi: 10.1063/1.1724339
- Paszke, A., Gross, S., Massa, F., Lerer, A., Bradbury, J., Chanan, G., et al. (2019). “PyTorch: an imperative style, high-performance deep learning library,” in *Advances in Neural Information Processing Systems* (Vancouver, BC), 8024–8035.
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830. doi: 10.5555/1953048.2078195
- Pierrard, V., and Lazar, M. (2010). Kappa distributions: theory and applications in space plasmas. *Sol. Phys.* 267, 153–174. doi: 10.1007/s11207-010-9640-2
- Rea, A., and Rea, W. (2016). How Many Components Should be Retained From a Multivariate Time Series PCA? *arXiv [Preprint]* arXiv:1610.03588
- Richardson, I. G. (2004). Energetic particles and corotating interaction regions in the solar wind. *Space Sci. Rev.* 111, 267–376. doi: 10.1023/B:SPAC.0000032689.52830.3e
- Richardson, I. G., and Cane, H. V. (2010). Near-earth interplanetary coronal mass ejections during solar cycle 23 (1996–2009): catalog and summary of properties. *Sol. Phys.* 264, 189–237. doi: 10.1007/s11207-010-9568-6
- Richardson, I. G., and Cane, H. V. (2012). Near-earth solar wind flows and related geomagnetic activity during more than four solar cycles (1963–2011). *J. Space Weather Space Clim.* 2. doi: 10.1051/swsc/2012003
- Richardson, I. G., Cliver, E. W., and Cane, H. V. (2000). Sources of geomagnetic activity over the solar cycle: relative importance of coronal mass ejections, high-speed streams, and slow solar wind. *J. Geophys. Res. Space Phys.* 105, 18203–18213. doi: 10.1029/1999JA000400
- Roberts, D. A., Karimabadi, H., Sipes, T., Ko, Y.-K., and Lepri, S. (2020). Objectively determining states of the solar wind using machine learning. *Astrophys. J.* 889:153. doi: 10.3847/1538-4357/ab5a7a
- Rougier, N., and Boniface, Y. (2011). Dynamic self-organising map. *Neurocomputing* 74, 1840–1847. doi: 10.1016/j.neucom.2010.06.034
- Sabine, E. (1852). VIII. On periodical laws discoverable in the mean effects of the larger magnetic disturbance—No. II. *Philos. Trans. R. Soc. Lond.* 142, 103–124. doi: 10.1098/rstl.1852.0009
- Satopaa, V., Albrecht, J., Irwin, D., and Raghavan, B. (2011). “Finding a “Kneedle” in a Haystack: detecting knee points in system behavior,” in *2011 31st International Conference on Distributed Computing Systems Workshops* (Minneapolis, MN), 166–171. doi: 10.1109/ICDCSW.2011.20
- Schwenn, R. (1983). *The Average Solar Wind in the Inner Heliosphere: Structures and Slow Variations* (Washington, DC: NASA conference publication).
- Schwenn, R., and Marsch, E. (eds.). (1990). *Physics of the Inner Heliosphere I: Large-Scale Phenomena*. Vol. 20. Springer Science & Business Media.
- Shlens, J., View, M., and Introduction, I. (2014). *A Tutorial on Principal Component Analysis*.
- Smith, C. W., L’Heureux, J., Ness, N. F., Acuña, M. H., Burlaga, L. F., and Scheifele, J. (1998). *The Ace Magnetic Fields Experiment BT-The Advanced Composition Explorer Mission*. Dordrecht: Springer Netherlands.
- Smith, E. J. (2001). The heliospheric current sheet. *J. Geophys. Res. Space Phys.* 106, 15819–15831. doi: 10.1029/2000JA000120
- Stakhiv, M., Landi, E., Lepri, S. T., Oran, R., and Zurbuchen, T. H. (2015). On the origin of mid-latitude fast wind: challenging the two-state solar wind paradigm. *Astrophys. J.* 801:100. doi: 10.1088/0004-637X/801/2/100
- Stakhiv, M., Lepri, S. T., Landi, E., Tracy, P., and Zurbuchen, T. H. (2016). On solar wind origin and acceleration: measurements from ACE. *Astrophys. J.* 829:117. doi: 10.3847/0004-637X/829/2/117
- Tsurutani, B. T., Gonzalez, W. D., Gonzalez, A. L. C., Guarnieri, F. L., Gopalswamy, N., Grande, M., et al. (2006). Corotating solar wind streams and recurrent geomagnetic activity: a review. *J. Geophys. Res. Space Phys.* 111. doi: 10.1029/2005JA011273
- Vettigli, G. (2013). *MiniSom: Minimalistic and Numpy Based Implementation of the Self Organizing Maps*.
- Villmann, T., and Claussen, J. C. (2006). Magnification control in self-organizing maps and neural gas. *Neural Comput.* 18, 446–469. doi: 10.1162/089976606775093918
- von Steiger, R., and Zurbuchen, T. H. (2015). Solar metallicity derived from *in situ* solar wind composition. *Astrophys. J.* 816:13. doi: 10.3847/0004-637X/816/1/13
- Wenzel, K. P., Marsden, R. G., Page, D. E., and Smith, E. J. (1992). The ULYSSES mission. *Astron. Astrophys. Suppl. Ser.* 92:207.
- Winterhalter, D., Smith, E. J., Burton, M. E., Murphy, N., and McComas, D. J. (1994). The heliospheric plasma sheet. *J. Geophys. Res. Space Phys.* 99, 6667–6680. doi: 10.1029/93JA03481

- Withbroe, G. L. (1986). *Origins of the Solar Wind in the Corona BT-The Sun and the Heliosphere in Three Dimensions*. Dordrecht: Springer Netherlands.
- Wu, C. C., Lepping, R. P., and Gopalswamy, N. (2006). Relationships among magnetic clouds, CMES, and geomagnetic storms. *Sol. Phys.* 239, 449–460. doi: 10.1007/s11207-006-0037-1
- Xu, F., and Borovsky, J. E. (2015). A new four-plasma categorization scheme for the solar wind. *J. Geophys. Res. Space Phys.* 120, 70–100. doi: 10.1002/2014JA020412
- Yin, H., and Allinson, N. M. (1995). On the distribution and convergence of feature space in self-organizing maps. *Neural Comput.* 7, 1178–1187. doi: 10.1162/neco.1995.7.6.1178
- Zhao, L., Landi, E., Lepri, S. T., Gilbert, J. A., Zurbuchen, T. H., Fisk, L. A., et al. (2017). On the relation between the *in situ* properties and the coronal sources of the solar wind. *Astrophys. J.* 846:135. doi: 10.3847/1538-4357/aa850c
- Zhao, L., Zurbuchen, T. H., and Fisk, L. A. (2009). Global distribution of the solar wind during solar cycle 23: ACE observations. *Geophys. Res. Lett.* 36, 1–4. doi: 10.1029/2009GL039181
- Zhao, L. L., Adhikari, L., Zank, G. P., Hu, Q., and Feng, X. S. (2018). Analytical investigation of turbulence quantities and cosmic ray mean free paths from 1995–2017. *J. Phys. Conf. Ser.* 1100:012029. doi: 10.1088/1742-6596/1100/1/012029

Conflict of Interest: The authors declare that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Copyright © 2020 Amaya, Dupuis, Innocenti and Lapenta. This is an open-access article distributed under the terms of the Creative Commons Attribution License (CC BY). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.