







Electromagnetic Ion Cyclotron Waves Pattern Recognition Based on a Deep Learning Technique: Bag-of-Features Algorithm Applied to Spectrograms

Claudia Medeiros¹ , V. M. Souza¹ , L. E. A. Vieira¹, D. G. Sibeck², B. Remya³, L. A. Da Silva^{1,4} , L. R. Alves¹,
J. P. Marchezi¹, P. R. Jauer^{1,4} , M. Rockenbach¹, A. Dal Lago¹, and C. A. Kletzing⁵

¹ Instituto Nacional de Pesquisas Espaciais 1758 Astronautas Av., Jardim da Granja São José dos Campos, SP 12227010, Brazil; claudia.medeiros@inpe.br, claudia.inpe@gmail.com

² NASA Goddard Space Flight Center, Greenbelt, MD, USA

³ Indian Institute of Geomagnetism, Navi Mumbai, Maharashtra, India

⁴ State Key Laboratory of Space Weather, National Space Science Center, Chinese Academy of Sciences, Zhongguancun Nanertiao, Haidian District, Beijing, People's Republic of China

⁵ Department of Physics and Astronomy University of Iowa, Iowa City, IA, USA

Received 2020 March 4; revised 2020 May 23; accepted 2020 May 26; published 2020 July 13

Abstract

Several studies have shown the importance of electromagnetic ion cyclotron (EMIC) waves to the pitch angle scattering of energetic particles in the radiation belt, especially relativistic electrons, thus contributing to their net loss from the outer radiation belt to the upper atmosphere. The huge amount of data collected thus far provides us with the opportunity to use a deep learning technique referred to as the Bag-of-Features (BoF). When applied to images of magnetic field spectrograms in the frequency range of EMIC waves, the BoF allows us to distinguish, in a semi-automated way, several patterns in these spectrograms that can be relevant to describe physical aspects of EMIC waves. Each spectrogram image provided as an input to the BoF corresponds to the windowed Fourier transform of a ~ 40 minutes to 1 hour interval of Van Allen Probes' high time-resolution vector magnetic field observations. Our data set spans the 2012 September 8 to 2016 December 31 period and is at geocentric distances larger than 3 Earth radii. A total of 66,204 spectrogram images are acquired in this interval, and about 45% of them, i.e., 30,190 images, are visually inspected to validate the BoF technique. The BoF's performance in identifying spectrograms with likely EMIC wave signatures is comparable to the visual inspection method, with the enormous advantage that the BoF technique greatly expedites the analysis by accomplishing the task in just a few minutes.

Unified Astronomy Thesaurus concepts: Planetary magnetosphere (997); Van Allen radiation belt (1758); Neural networks (1933); Space plasmas (1544); pp waves (1288); Support vector machine (1936); Classification (1907)

1. Introduction

The Earth's magnetosphere, which is formed by the interaction of the magnetized solar wind with Earth's magnetic field, contains a region known as the Van Allen radiation belts. In this region, energetic particles are mostly trapped in the inner belt between 1.1 and 2 Earth radii ($1 R_E = 6371$ km) and in the outer belt approximately between 3 and $6 R_E$. Particle fluxes in the inner belt generally do not vary appreciably over the course of months or even years. In the outer radiation belt, however, the story is quite different. An increase in solar wind dynamic pressure can reduce the radiation belt electron fluxes. In such a way, the filling in of the radiation belt is controlled mainly by the solar wind velocity (see e.g., Li et al. 2005; Reeves et al. 2011; Wing et al. 2016, and references therein). In addition, these fluxes can be altered in a matter of hours by both external (to the Earth's magnetosphere) and internal causes, such as magnetopause boundary motions driven by solar wind dynamic pressure fluctuations (see e.g., Souza et al. 2017, and references therein), relativistic electron gyro-resonant interactions with internally generated electromagnetic waves (see e.g., Shprits et al. 2009; Usanova et al. 2014, and references therein), and other wave particle interactions (see more details in, e.g., Roederer & Zhang 2016; Jaynes & Usanova 2020).

Particle fluxes in the outer radiation belt can change by orders of magnitude either over periods of 12 hr or even less (Roederer & Zhang 2016; Jaynes & Usanova 2020) or remain essentially unaltered by many days (e.g., Alves et al. 2016;

Ozeke et al. 2017). Such rapid variations in the outer belt fluxes with subsequent removal of particles from the outer radiation belt have been attributed to the gyro-resonant interaction with the so-called electromagnetic ion cyclotron (EMIC) waves whose frequencies typically range from 0.1 to 5 Hz (see e.g., Thorne & Horne 1992; Summers & Thorne 2003; Zhang et al. 2010; Thorne et al. 2013; Usanova et al. 2014; Medeiros et al. 2019). Studying the characteristics of such interactions, as well as EMIC wave occurrence patterns in the radiation belts region, is extremely important since EMIC waves contribute significantly to the variability of the energetic particle fluxes in the outer radiation belt and also promote the precipitation of radiation belt particles into the upper atmosphere, thereby affecting its chemical composition (Horne et al. 2009). Space missions such as NASA's Van Allen Probes (Mauk et al. 2012) provide us a data set of in situ measurements made by electric and magnetic field sensors through which we can infer EMIC wave occurrence (Kletzing et al. 2013).

EMIC waves identification in near-Earth space has traditionally been performed via Fourier analyses of locally measured magnetic field data. The Fourier power spectral density (PSD) in units of $nT^2 \text{ Hz}^{-1}$ is plotted in a frequency versus time graph, yielding the so-called spectrograms. EMIC wave detection in such spectrograms is made: (i) by eye inspection of the wave packets' magnetic field amplitude, which in turn translates to (usually) localized, in both time and frequency, enhancements of PSDs relative to background

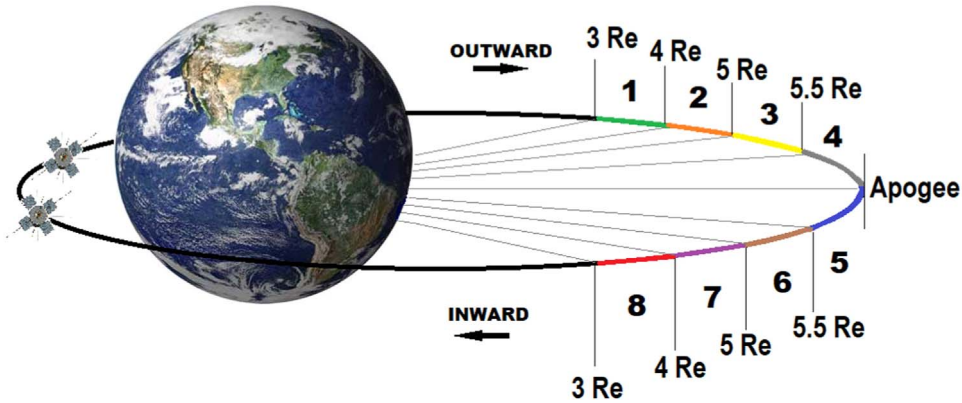


Figure 1. Sketch of regions along the Van Allen Probes orbit where the stacked up spectrogram images (see Figure 2) are obtained. These numbered regions (from 1 to 8) are selected according to the satellites’ radial distance in Earth radii. The time it takes for the satellites to cross each region is around 50 minutes.

values or (ii) by computational techniques that interactively analyze the pixel intensities of the spectrogram’s images against the background (see e.g., Usanova et al. 2014; Saikin et al. 2015; Wang et al. 2015, and references therein). When wave packets are clearly distinct from the background, the interval can be considered as having a candidate EMIC wave event. Since the waves are in the 0.1–5 Hz frequency range, the data acquisition rate must be at least twice as high to resolve the wave packets. The Van Allen Probes mission, for example, provides local magnetic field measurements at an acquisition rate of 64 Hz, i.e., approximately one sample every 15 ms, which is high enough to detect EMIC waves. In the case where only a short period needs to be investigated, visual inspection seems to suffice. However, when the data set span months or years, individually checking EMIC wave intervals that typically range from a few minutes to a few hours may be impractical. This study aims to present a new technique referred to as the Bag-of-Features (BoF) that is based on deep learning algorithms. It quickly and efficiently finds signatures of EMIC waves in large amounts of spectrogram images. Data from NASA’s Van Allen Probes mission are used here as an application example, but the technique is general enough so it can also be applied to any kind of spectrogram images. This paper is organized as follows: Section 2 presents details of the Van Allen Probes’ data used here, whereas Section 3 provides a brief description of EMIC waves. Next, in Section 4, a brief introduction to artificial neural networks (ANN) is presented as well as the deep learning-based technique used for both identification of salient features in spectrogram images and their classification into user-defined classes. After that, Section 5 presents the methods employed here—specifically how the spectrograms and the corresponding images are generated—and the procedures used for validating the proposed technique. Finally, Section 6 summarizes the results obtained when the BoF technique is applied in the identification of likely EMIC wave events in a Van Allen Probes data set spanning a ~ 4 yr period.

2. Data Set Description

The Van Allen Probes mission consists of two low inclination (10°), elliptically orbiting, and identically instrumented satellites named the Radiation Belt Storm Probes (RBSP-A and RBSP-B), which were launched on 2012 August 30. Both perform approximately three orbits in a daily basis with an orbital period of approximately 9 hr. The spacecraft altitude ranges from $0.1 R_E$ to $5.8 R_E$ at apogee and the satellites lap each other during the

orbit with lags ranging from a few minutes to 4.5 hr (Mauk et al. 2012). They cross the inner and outer Van Allen radiation belts, measuring a number of physical parameters. Our main interest is the high time-resolution magnetic field data obtained by the Electric and Magnetic Field Instrument Suite and Integrated Science (EMFISIS) instrument (Kletzing et al. 2013). The EMFISIS magnetometer on board both satellites measures the magnetic field vector with a sampling rate of 64 Hz. The data set were available in several coordinate systems. In this study, we choose data from all three magnetic field components B_x , B_y , and B_z in Geocentric Solar Magnetospheric (GSM) coordinate system, which has its X -axis along the Earth–Sun line, its Y -axis perpendicular to the Earth’s magnetic dipole, and its X – Z plane in the dipole axis. The data set interval in this study comprises the period from 2012 September 8 to 2016 December 31. We limit our data set to measurements acquired mostly in the outer Van Allen belt, i.e., at altitudes above $3 R_E$. The radiation belt community typically uses the L^* parameter to address radial distance from Earth. The L^* parameter was not relevant to perform the PSD calculation itself but it could increase significantly the computer-time demand. Thus, we adopted radial distance obtained directly from the high-resolution GSM data set. The orbital period is divided into eight segments according to the spacecraft radial distance, as shown in Figure 1. This is done so we can discriminate EMIC waves’ occurrence at different radial distances. In this way, the orbital period is split in regions in each of which the spacecraft spent 40–60 minutes. Magnetic field vector components are stored in each of these time windows, and they are subsequently Fourier analyzed, as described below in Section 5.1, in order to generate the spectrogram images that will be used as an input to the neural network-based tool presented here.

3. EMIC Waves

In the context of Earth’s Van Allen radiation belts, EMIC waves play a crucial role on the variability of the relativistic ($\gtrsim 1$ MeV) electron population (Reeves et al. 2003; Horne et al. 2009; Li et al. 2014; Usanova et al. 2014; Zhang et al. 2016b; Clilverd et al. 2017; Jaynes & Usanova 2020). They are commonly found in the 0.1–5 Hz frequency range. These waves are associated with ion cyclotron instabilities, which are usually excited by ring current ion injections during geomagnetic storms and also during compressions of the day-side magnetopause: the boundary that separates the Earth’s magnetized plasma from the solar wind (Cornwall 1965; Kennel & Petschek 1966; Summers & Thorne 2003; Thorne et al. 2013). The local plasma composition defines the EMIC wave’s



Figure 2. PSDs in units of $\text{nT}^2 \text{Hz}^{-1}$, obtained from a Fourier analysis of the time series of three magnetic field components measured by the EMFISIS instrument on 2015 April 16 from 05:33 to 06:10 UT. The y-axis ranges from 0 to 5 Hz on the left-hand side panels, whereas it goes from 0.2 to 5 Hz on the right-hand side image. Panels (a), (b), and (c) contain the PSDs of δB_x , δB_y , and δB_z , respectively, along with the gyrofrequencies of ions of hydrogen (green line), helium (red), and oxygen (black line). These panels show a typical EMIC wave signature found in our data set. On the right-hand side, all three PSDs were stacked up to form a single image (560×420 pixels) containing only the PSD information, so it can be used as an input to the BoF technique. See the text for details.

propagation frequency band and the energy level of the particles that will be affected by their occurrence. Ions of hydrogen (H^+), helium (He^+), and oxygen (O^+) are present in the outer radiation belt and their gyrofrequencies define the bands of propagation of these waves. The most common propagation band is the hydrogen band, as shown in Figures 2(a)–(c), which presents a typical EMIC wave signature found in our data set. A burst of intensified PSD values reaching around $\sim 10^0$ – $10^1 \text{ nT}^2 \text{Hz}^{-1}$ can be seen between the gyrofrequencies of hydrogen (green line) and helium (red line). The waves are mostly left-hand polarized and their amplitude can range from a few to tens of nT (Clausen et al. 2011; Halford et al. 2016). The preferential region of occurrence of EMIC waves is known to be the afternoon magnetic local time (MLT) sector from $\sim 12:00$ to $\sim 18:00$ MLT in the region near the plasmapause and the plasmaspheric plume (see e.g., Saikin et al. 2015; Wang et al. 2015; Tetrick et al. 2017, and references therein). The occurrence of EMIC waves and the resonant interaction with energetic particles promotes pitch angle scattering and the likely precipitation of these particles into the upper atmosphere (Li et al. 2014; Usanova et al. 2014; Remya et al. 2015; Zhang et al. 2016a; Clilverd et al. 2017).

The most common techniques to identify EMIC waves is the analysis of their PSDs in order to find signatures of their occurrence, intensity, and frequency band propagation (see e.g., Usanova et al. 2014; Saikin et al. 2015; Wang et al. 2015; Medeiros et al. 2019, and references therein). Afterward, these data are validated by detailed analysis of the electric and/or magnetic fields as a function of direction of propagation, wave mode, and polarization. However, visual inspection to find EMIC waves signatures is quite time consuming, especially when it is necessary to sweep through a large spectrogram database. Thus, it is highly desirable to use other automated search methods, particularly those based on deep learning techniques, that significantly narrow down the number of likely EMIC wave event candidates.

4. Neural Network-based Technique: The BoF Approach

ANN have been gaining an ever increasing space of applicability in many disciplines, particularly in space sciences (Camporeale et al. 2018). This success is due mostly to the significantly large amount of accumulated data on which ANN are especially well suited to be applied. Some of the ANN’s goals include the identification and clustering of patterns by using computational systems in a faster way. An ANN is essentially made up of parallel-working simple processing units known as “neurons,” which have the “propensity for storing experiential knowledge and making it available for use” (Haykin 2008). This empirical knowledge is gained by the ANN via several exposures to a given data set, and it is stored in the form of synaptic weights that essentially encode the “importance” of the input data to each neuron pertaining to the ANN. By devising an algorithm intended to iteratively modify the synaptic weights, one can “train” an ANN to perform a specific task, which can be, but is not limited to, data clustering and classification, pattern recognition, image and data compression, and so forth. For an in-depth view on the inner workings and applications of ANNs, the reader is referred to Haykin (2008).

There is a very important aspect of the learning/training phase of an ANN that is explored in this work, and it has to do with *how* the learning process is carried out, i.e., whether it is supervised or not. In a supervised training one already knows the expected outputs the ANN should provide in advance, so the ANN structure is modified accordingly. In this case, the user effectively interferes in the ANNs’ training phase. On the other hand, in an unsupervised training, the ANN itself will try to create its own organization or representation of the information it receives during learning time without any interference from the user.

Section 5.3 describes the so-called BoF technique (O’Hara & Draper 2011), which is a deep learning technique that employs both of the learning paradigms mentioned above. By using a

large database of spectrogram images as an input to the BoF, our goal is to recognize patterns in those images that are characteristic of EMIC wave signatures. Briefly, the different patterns, or *features*, are used to group the images into distinct, user-defined classes each containing a set of characteristics of the input data. More details on the feature extraction algorithm employed here and classification/grouping of images can be found in Sections 5.4.1 and 5.4.2, respectively. The main advantage of using the BoF technique is that it is based on a methodology that allows for a more efficient processing of the image structure (see Vapnik 1999, for details). Moreover, it is a simple algorithm to implement, using few computational resources and accomplishing the task typically in a few minutes that otherwise would take days or even weeks to complete.

We used a BoF implementation provided by the Imaging Processing, Statistics and Machine Learning Toolboxes from MATLAB™. In what follows, we describe in detail all the steps employed to implement the BoF technique.

5. Methodology

5.1. PSD Evaluation and Preparation of Input Data

The PSD is obtained via a Fourier analysis of the magnetic field component perturbation $\delta B_i = B_i - B_{\text{avg},i}$, where $i = x, y$, and z and $B_{\text{avg},i}$ is a 100 s time average of the magnetic field component B_i , similar to the method employed by Medeiros et al. (2019). Specifically, we performed a short-time Fourier transform (STFT) of the δB_i time series with a window length of 3750 samples (~ 1 minute) and an overlap of 1875 samples (~ 30 s). An example of spectrograms generated using the STFT method is shown in Figures 2(a)–(c). The frequency (vertical) axis range is set from 0 to 5 Hz. Later on, all three magnetic field components have their spectrograms stacked up as shown in the right panel of Figure 2 in order to form a single 560×420 pixels image. Any white spaces that surround the resulting stacked up spectrogram image are trimmed off so that the area within the image being analyzed by the BoF technique contains only the spectrograms. Failing to perform this step might result in a poor BoF performance. Such an edited image is used as an input to the BoF. Notice that in the stacking up process, we deliberately increased the lower bound of the frequency axis range to 0.2 Hz to exclude the higher PSD values contained in such lower frequencies. They were posing a problem for the feature extraction algorithm used within the BoF (see Section 5.4.1 for details) in the sense that most of the features found were located exactly over this lower frequency range, i.e., from 0 to 0.2 Hz, and as a result, a number of likely EMIC wave events were being overestimated. All images analyzed here are generated in the same way, and they all have the same color bar range—namely, from 10^{-4} to 10^2 nT² Hz⁻¹.

5.2. Separating Input Data into Classes

Most of the time during the Van Allen Probes orbits, one does not find EMIC wave signatures in the spectrograms, thus a large number of unwanted cases are also present in our data set amidst the sought-after EMIC wave events. Therefore, we deemed it wise to allow the BoF technique to group together the stacked up spectrogram images into distinct classes, where each class would hold within it images that share some characteristics. The number of classes wherein the data set is going to be divided is defined by the user. Visual inspection of about 10% of our data set (that is,

~ 6600 images) revealed spectrograms often showing: (i) no discernible PSD enhancement above the background level, like those shown in Figure 3(c); (ii) instrumental artifacts, as shown in Figure 3(d); (iii) clear EMIC wave signatures, as presented in Figures 2 and 3(b); and finally (iv), a mixture of instrumental artifacts and PSD enhancements that could possibly be characteristic of EMIC wave signatures but are not as clear as those seen in the images categorized as (iii), thus requiring a careful analysis to find unambiguous evidences of EMIC waves. This way, it seemed natural to divide the data set into 4 classes. However, a preliminary analysis of the BoF technique’s performance has shown that many of the images that by visual inspection would fall into categories (i), (ii), and (iv) were erroneously categorized as EMIC cases, which considerably decreased the BoF accuracy (more on that in Section 5.4.2). In order to minimize misclassification of EMIC cases by the BoF technique, we then chose to group images that were visually categorized as (iii) and (iv) as forming a new class referred to as EMIC, and all other types of spectrograms were gathered into a single class named NOEMIC.

Our data set is comprised of 66,204 stacked up images, all with the same format as the one shown on the right-hand side of Figure 2. Data from the first, second, and half of the third precession periods analyzed in more detail in Section 6.3 correspond to the following time intervals, respectively: 2012 September 8 to 2014 June 19, 2014 June 20 to 2016 March 29, and 2016 March 30 to 2016 December 31. About 45% of the total (that is, 30,190 images) are visually inspected and classified into either one of the two major classes above: EMIC or NOEMIC. This step, although tedious, is very important for the validation of the BoF technique, as shown below in Section 5.4.3. These 30,190 images spanned a period going from 2012 September 8 to 2014 August 25 that accounts for somewhat more than the first full orbit precession of the Van Allen Probes around Earth (it actually completed the first precession about two months earlier, on June 19), so the orbit’s apogee covered all MLTs.

5.3. BoF’s Training and Validation Data Sets

As with any other deep learning technique, the BoF needs to be fed with some subset of the input data to be trained. The BoF, and other ANN-based techniques, has the powerful ability to generalize, i.e., it can provide reasonable outputs for inputs not encountered during the training phase (Haykin 2008). Thus, to train the BoF, one does not need to use the whole data set but rather only a smaller, representative part of it. As mentioned in the previous section, the total amount of visually classified stacked up spectrogram images are 30,190, and they are obtained within the 2012 September 8 to 2014 August 25 period, as already mentioned above. Such images were not randomly chosen. Instead, we inspected each and every one of those 30,190 images in the same sequence they were acquired. This data set is nearly evenly distributed between the two satellites, i.e., each contributed with about 15,000 images. These images are then split into 2000 images for composing the training data set and 28,190 for the validation one. Smaller numbers for the training data set have also been tested, but the 2000 image set was found to be the one that provided a slightly better accuracy. By “accuracy” we mean the percentage of correct classifications performed by the BoF, i.e., we check whether or not the BoF’s classification of the input images matches the previously done visual classification of these same input images. The following sections are devoted to provide

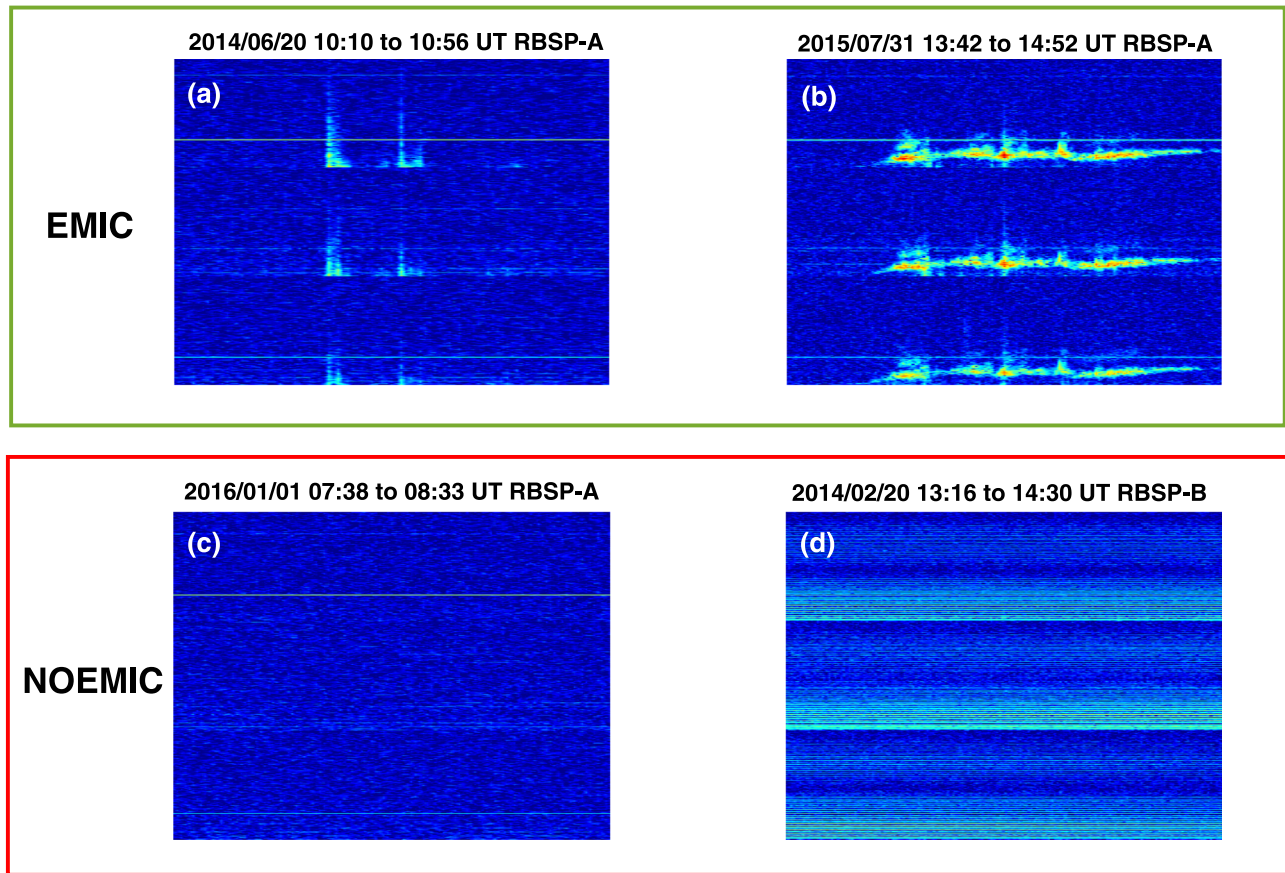


Figure 3. Typical stacked up spectrogram images found in our data set. The date, time interval, and which spacecraft the spectrograms have been obtained with are shown on top of each panel. Images like those shown on panels (a) and (b) form a single class, named EMIC, with likely EMIC wave events. They exemplify cases where a PSD enhancements might be related with EMIC wave signatures, but a further analysis is required, and (b) where unambiguous evidence of EMIC waves is found. Likewise, images like those shown on panels (c) and (d) form a single class, named NOEMIC, where (c) no distinguishable EMIC wave signatures are present and (d) there is a presence of instrumental artifacts that do not allow a clear eye inspection of EMIC wave signatures.

more details on the BoF’s classification process and accuracy determination.

The training data set possesses 1000 images randomly taken from the visually classified EMIC class and another 1000 images also randomly taken from the visually classified NOEMIC class. Both Van Allen Probes contributed about 50% of the number of images in each class, as summarized in Table 1, and the random selection ensured that both EMIC and NOEMIC cases would be fairly well distributed in MLT. The other 36,014 images acquired during the time interval from 2014 August 26 to 2016 December 31 have *not* been visually classified, and they also serve to be used as an input to the BoF and to be classified by it, as shown in Section 6.

5.4. Applying BoF for Identification and Classification

The stacked up images are generated as described in Section 5.2. The steps below detail how the BoF classifies them. The entire process is divided in two blocks: (i) training process and accuracy evaluation and (ii) application to unclassified spectrogram images, as shown in the schematic diagram in Figure 4. Consider initially the upper block. After visually classifying 30,190 images, 2000 of them will compose the so-called “training data set,” as already pointed out in the previous section. Such images are then used as an input to the BoF, which is represented in Figure 4 by the dashed box.

Table 1
Classification of Stacked Up Spectrogram Images as Either EMIC or NOEMIC

Data Set	EMIC		NOEMIC		Total
	A	B	A	B	
Training	538	462	466	534	2000
Validation	1416	1074	12,674	13,026	28,190
BoF classification	3920	3787	14,371	13,936	36,014
Total	5874	5323	27,511	27,496	66,204

Note. Images pertaining to both training and validation data sets totaling 30,190 images are classified via visual inspection, whereas the remaining 36,014 images from the whole data set are classified only via the BoF technique. A and B denote RBSP-A and RBSP-B, respectively.

Notice that the BoF technique is composed of two other intermediate steps referred to as “feature extraction” and “K-means clustering,” which are presented in the next sections. Afterward, the images composing the validation data set are classified by the BoF, and finally, the BoF’s accuracy can be evaluated. Subsequently, in the bottom block, the BoF is used to classify the remainder of our data set, i.e., 36,014 images, whose results are presented in Section 6.

The following subsections provide more details of the BoF technique and its application on the classification of stacked up spectrogram images.

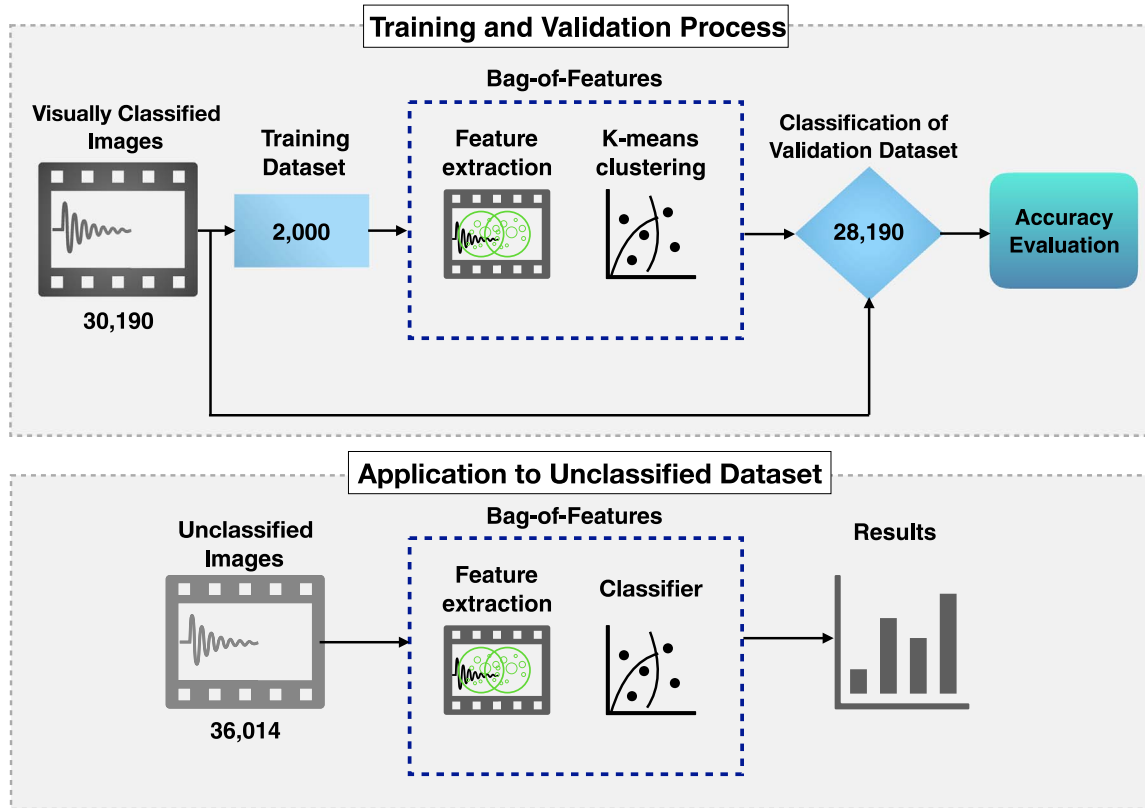


Figure 4. Schematic diagram to illustrate all steps involved in the classification of input images. A total of 30,190 visually classified images are used to compose both the training (2000 images) and validation (28,190 images) data sets. Once presented to the BoF, the images have their features extracted and an image classifier is built via K-means clustering (see the text for details). Using the validation data set, the BoF’s accuracy can be evaluated. Afterward, the unclassified (36,014) images are presented to the BoF to be classified and the results are further analyzed.

5.4.1. Feature Extraction

The main target here is to find regions within the image, hereafter named *features*, with relevant information so one can use them to distinguish one image from the other. In our data set, this means to find features in spectrogram images that can be identified as EMIC wave signatures. It turns out that such features are accompanied by other types of features that characterize regions in an image that are *not* related with EMIC wave signatures but are related to background noise and instrumental artifacts instead. Thus, the challenge here is to find an optimum number of features so one can use them to identify images possessing similar types of features and then grouping such images into classes. These classes will be characterized by images that hold similar characteristics.

Features are found using the Speeded Up Robust Features (SURF) detector (see Bay et al. 2008, for details), and here we describe its main characteristics. SURF’s goal is to find a certain amount of features in an image and to chart each one of them into a “vector descriptor,” which is a mathematical way of unambiguously identifying a feature in the image. These concepts are made clear below.

The first step to find *interesting* features in the image is having a way to detect the so-called *interest points*. These are usually locations in the image where there are strong spatial variations, i.e., gradients, in the pixel intensities relative to the points surrounding the interesting features. To detect them, we first need to convert our red, green, and blue color code, stacked up spectrogram images to a gray scale, since that is how the SURF detector operates. The pixel intensities varied from 0 (black) to

255 (white), and they are denoted by $I(i, j)$, where integer indexes i and j refer to a given pixel location in the image. After that, an important procedure is performed on the input image I that accounts for SURF’s reduced computational time: the *integral image* representation I_{Σ} of the gray-scale image I . An element of an integral image $I_{\Sigma}(\mathbf{r})$ at a location $\mathbf{r} = (x, y)$ represents the sum of all pixels in the input image I of a rectangular region formed by the point \mathbf{r} and the origin, that is

$$I_{\Sigma}(\mathbf{r}) = \sum_{i=0}^x \sum_{j=0}^y I(i, j). \quad (1)$$

The reason for using the integral image representation is as follows. SURF is a Hessian-based detector (detailed below) and it recurrently performs convolutions of the image I with so-called box-type filters, which are very basic approximations for Gaussian second order derivatives (see Figure 1 of Bay et al. 2008, for details on the box filters). Therefore, a large number of addition operations are needed, particularly when the image size, i.e., the number of pixels, is big. Once I_{Σ} is calculated, the number of additions to be done as a result of the convolution operations is drastically reduced, especially when the aforementioned box-type convolution filters are used.

SURF uses the Hessian matrix formulation to detect pixel intensity variations in the image along the horizontal, vertical, and diagonal directions. This is a required step to find the interest points whose determination is based on local maximization of the determinant of the Hessian matrix in the vicinity of a given pixel in the image I . Considering a point $\mathbf{r} = (x, y)$ in input image I , the Hessian matrix $\mathcal{H}(\mathbf{r}, \sigma)$ in \mathbf{r} at a

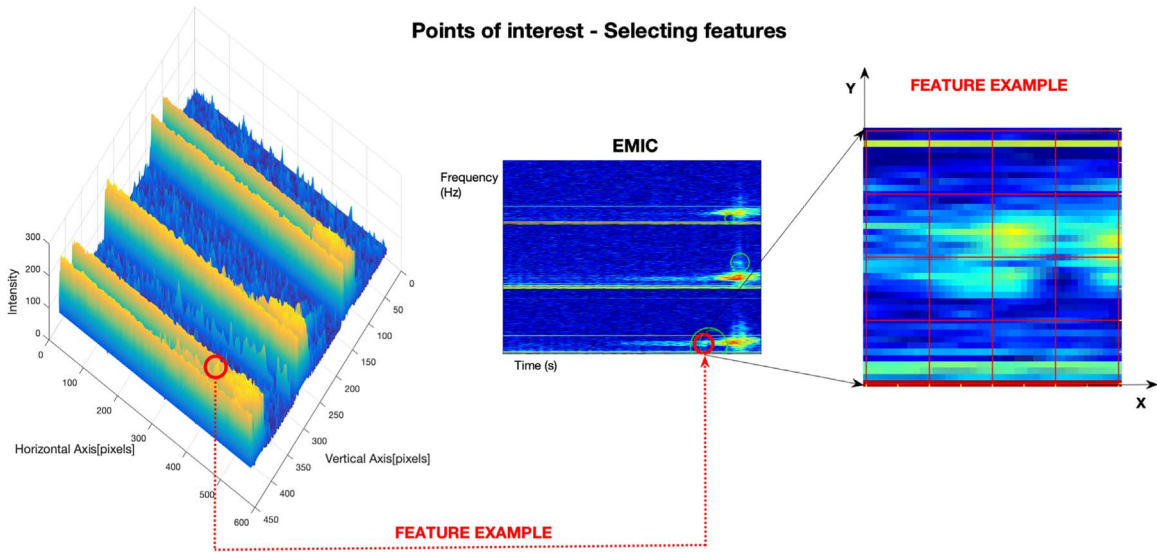


Figure 5. An example of feature selection by the SURF algorithm. (Left): three-dimensional representation of an example stacked up spectrogram image with the vertical and horizontal axis providing an estimate of the image’s pixel size. The third, out-of-plane axis represents the pixel intensity. (Middle): two-dimensional representation of the image on the left showing a typical EMIC wave signature. The center of the green circles represent the interest points’ locations in the image as determined via the SURF algorithm (see the text for details). Tens of thousands of features are typically extracted, but only five of them are shown for clarity. (Right): a blown up image feature. The feature is split up into 4×4 subregions, with each subregion having a four-dimensional vector descriptor containing the underlying intensity structure of that subregion. Ultimately, each feature is unambiguously represented by a more general vector descriptor with 64 elements.

scale σ is given by

$$\mathcal{H}(\mathbf{r}, \sigma) = \begin{bmatrix} D_{xx}(\mathbf{r}, \sigma) & D_{xy}(\mathbf{r}, \sigma) \\ D_{xy}(\mathbf{r}, \sigma) & D_{yy}(\mathbf{r}, \sigma) \end{bmatrix}, \quad (2)$$

where $D_{xx}(\mathbf{r}, \sigma)$ is the convolution of the approximate second order Gaussian derivative with the image I at point \mathbf{r} and likewise for $D_{xy}(\mathbf{r}, \sigma)$ and $D_{yy}(\mathbf{r}, \sigma)$. The determinant of the Hessian matrix provides us with quantitative information about pixel intensity variation in the immediate neighborhood of a given pixel and at a given scale, σ . The scale parameter σ effectively enters as a smoothing parameter (that is, the higher its value, the higher the degree of smoothing of the original image). By smoothing the image, one increases the suppression of noise and other interfering fine-scale structures, so this procedure is expected to facilitate the search for interest points that uses maximization of the determinant of $\mathcal{H}(\mathbf{r}, \sigma)$ across different scales (see Bay et al. 2008, and references therein for more details on the interest points determination). When σ is increased, the box-type convolution filter’s size also increases, hence there is an upper limit to which σ can be increased, otherwise the box filter’s size would surpass the image size. For the stacked up spectrogram images used here, eight scale values are used that translate into eight box-type convolution filter sizes of 9×9 , 15×15 , 21×21 , 27×27 , 39×39 , 51×51 , 75×75 , and 99×99 pixels. Further increasing σ would only make the computational time for feature selection grow bigger without bringing newer information that has not been already found in the aforementioned eight scales (Bay et al. 2008).

Locating an interest point in the input image I is tantamount to say that its $(x, y; \sigma)$ location has been found. The interest point’s (x, y) location marks the centroid of the feature it is related to. The scale σ wherein the interest point has been located ultimately determines the size of the corresponding

feature in the image, i.e., the higher the σ value, the larger the feature’s size. The central panel in Figure 5 shows an example of the results of the feature extraction algorithm used in this study. In fact, a few thousands of features are usually extracted, but for clarity, only five of them are shown. They match locations in the image where one would normally associate the higher PSD values as being characteristic of an EMIC wave signature.

Each located feature has associated with it a vector descriptor that, as already mentioned above, is a mathematical way of unambiguously identifying that feature. Each feature, regardless of its (pixels) size, is represented by an unique array with 64 elements. The 4×4 square subregions shown on top of the highlighted feature on the right-hand side of Figure 5 just emphasizes the descriptor’s structure. Each subregion has a four-dimensional descriptor vector for its underlying intensity structure, thus resulting in a more general descriptor vector of 64 in length. Each element’s determination is described in great detail in Section 4 of Bay et al. (2008). The key point here is that by having such descriptors, one can perform a quantitative comparison among features pertaining to different images and then be able to group/classify images that have similar characteristics. The classification process is presented next.

5.4.2. Classification of Input Images

The classification process in this study is done in a supervised way, in the sense that the user specifies in advance the number of classes into which the BoF will group the input images. The main target of this step is to obtain what it is called here as the classifier. It starts with presenting the set of 2000 visually classified images used for the training/learning phase to the BoF. Recall that we already split up these images into two equally sized classes—namely, EMIC and NOEMIC—and provide them as inputs to the BoF in this way. Examples of such images are presented in the upper part of Figure 6. The feature extraction (SURF) algorithm, as briefly presented in

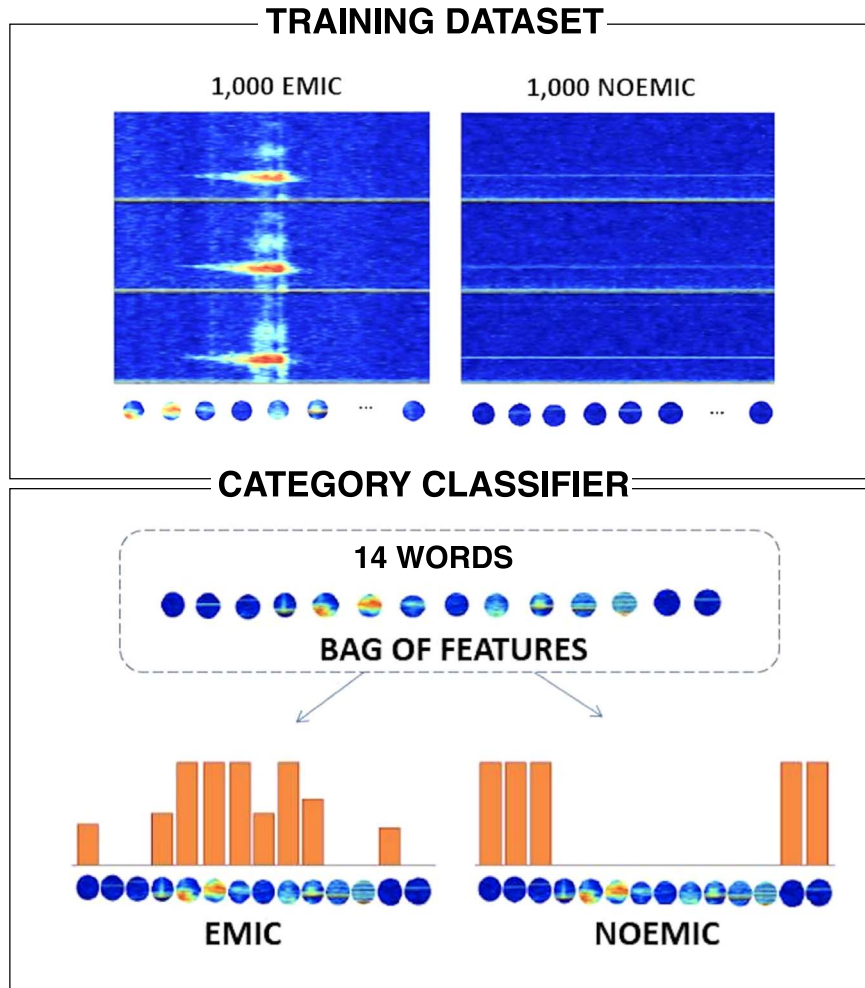


Figure 6. An illustration of the classification process. During the training phase, a set of 1000 images for each of the two classes, i.e., EMIC and NOEMIC, have their features extracted by the SURF detector. A K-means clustering procedure (see the text for details) is used to group the extracted features into an user-defined number of features, or *words*, which will compose the BoF. In this illustration, the “bag” contains only 14 words (features), while for the configuration used in the present study, 500 words are used. Next, a new (unclassified) image goes through the classifier where the image has its features extracted via SURF and they are subsequently compared against those in the BoF. Histograms like the ones shown on the bottom are constructed by finding the features extracted from the new (unclassified) image that is closest to those in the BoF (see Dietterich & Bakiri 1995, for details). Images classified as likely EMIC wave candidates typically show histograms like the one on the left, whereas images classified in the NOEMIC group present typical histograms like the one on the right.

the previous section, is applied to each and every image within the training data set, so we end up with a total number of extracted features, N_f (or vector descriptors), for each of the two predefined classes. The larger number of extracted features is always found in the EMIC class (that is, $N_{f,EMIC} > N_{f,NOEMIC}$), since it contains images possessing higher pixel contrasts relative to the background.

The numbers $N_{f,EMIC}$ and $N_{f,NOEMIC}$ are commonly on the order of tens of thousands, so what BoF does is select a number of features from each class to compose a single BoF containing an user-defined number of features, $N_{f,user}$, or *words*. The $N_{f,user}$ features (i.e., *words*) are presumably the most relevant ones that shall be subsequently used in the classification of both the validation and unclassified data sets. The feature selection process for composing the BoF is done via K-means clustering (Duda & Hart 1973) wherein the user defines the number $N_{f,user}$ of cluster centers whereby the vector descriptors of the training data set are going to be grouped. Then, this “bag” with $N_{f,user}$ features is the classifier. We tested several values of $N_{f,user}$, as shown in Table 2, and $N_{f,user} = 500$ provided both a slightly

better mean accuracy when compared to other smaller values and the best accuracy for the NOEMIC cases (it is also important to exclude cases that we are not interested in from our data set). Whenever a new image is presented to the BoF, its features are extracted with SURF and then compared against those in the “bag” and then a histogram of features’ occurrence like those shown on the bottom panel in Figure 6 are obtained (see Dietterich & Bakiri 1995, for details on how to compare extracted features with those in the “bag”). Notice that images classified as EMIC (NOEMIC) by the BoF technique have histograms that resemble the one being shown on the left (right).

After the classification process is done, one needs to verify its accuracy by presenting to the BoF the validation data set that has been already visually classified.

5.4.3. Validation Process

The validation data set contained 28,190 images from where 2490 are visually classified as EMIC and 25,700 are classified

Table 2

BoF's Performance Evaluation According to the Number of Words (See the Text for Details)

Words or $N_{f,user}$	Average Accuracy (%)	NOEMIC (%)	EMIC (%)
500	86	94	78
250	85	92	79
125	86	91	81
100	86	91	80
50	84	88	80
25	81	87	76
10	80	83	76

Note. Bold values have been selected to denote the word parameter $N_{f,user} = 500$ since it provides the best combined evaluation: for the classification of the NOEMIC cases (94%), EMIC cases (78%), and an average accuracy of 86%.

as NOEMIC. After being presented to the BoF's classifier, as schematically shown in the upper block of Figure 4, a so-called confusion matrix is generated comparing the number of times, in percentage, that the BoF's classifier correctly or incorrectly predicted the class wherein the input image should be categorized. The confusion matrix is presented in Table 3. The percentages shown in Table 3 are obtained as follows. All 2490 EMIC cases are exposed to the trained neural network and 78% of them (1931 images) are correctly classified into the EMIC class. The remaining 22% (559 images) are then computed as pertaining to the NOEMIC class, which in turn contributes to increase the error in the classification of NOEMIC-type images. Likewise, when the 25,700 NOEMIC cases are exposed to the trained neural network, 94% (24,171) of them are correctly classified into the NOEMIC class, and the remaining 6%, which account to 1529 images, fall directly into the EMIC class, which ends up overestimating the number of EMIC-type images.

A possible explanation for such a discrepancy is the following. The NOEMIC-type images contain mostly the (blue) background PSD levels, thus the extracted features from those images are not so different from one another. As a result, the BoF's classifier does not find major problems in correctly matching the extracted features from these NOEMIC-type images with those contained in the BoF. The same does *not* happen for the EMIC-type images, which have a very large number of distinct features, so the BoF's classifier is more prone to misclassifications in this case. Moreover, we have found that the BoF's classifier has some difficulty disentangling small EMIC waves' signatures from the background noise, as illustrated in Figure 7 where panels (a) and (b) present examples of images visually classified as belonging to the EMIC class and the BoF's classifier categorized them as pertaining to the NOEMIC class. Notice that EMIC waves' signatures are very faint in both panels (a) and (b), as indicated by the red ellipses, so it might be quite hard for the algorithm to automatically separate those from instrumental artifacts that are superposed on the image as horizontal green stripes. Another example of misclassification is shown in panels (c) and (d) in Figure 7, where the images are visually classified as an NOEMIC-type and the BoF's classifier categorize them as pertaining to the EMIC class. Both horizontal red stripes in panel (c) and the vertical, very faint stripes amidst the

Table 3
Confusion Matrix of the Validation Data Set

Known	Predicted	
	EMIC %	NOEMIC %
EMIC	78	22
NOEMIC	6	94
Average accuracy is 86%		

Note. Bold values stand for results that are correctly evaluated in 78% of the EMIC cases and in 94% of the NOEMIC cases after exposing the validation data set to BoF's classifier.

background in panel (d) most likely contributed for the misclassification. Still, the 86% average accuracy of the BoF's classifier is quite acceptable.

6. Results


6.1. BoF's Classification Performance Compared to Visual Classification

Before presenting the BoF's performance with the unclassified set of images, it is important to see how well the visually and BoF classified images agree, or not, with each other. That is, in order to be trusted, we expect the neural network-based technique to perform nearly as good as the human eye during the classification process or even to outperform the method of visual classification. Figure 8 tries to address this point. It shows the number of likely EMIC wave cases, per month, as a function of time, as obtained by visual classification (blue line) and via the BoF technique (red line). Only the validation data set has been used in this case, and it spans the 2012 September 8 to 2014 August 25 period. One can see that nearly throughout the analyzed period, the BoF technique performs quite well when compared to the visual classification method. It closely follows the blue curve, overestimating the monthly number of likely EMIC cases usually by a factor of less than 4, which we consider to be quite acceptable. There are two contributing factors to the overestimation in the EMIC-like cases: one has been already mentioned in Section 5.4.3, which has to do with how the BoF technique handles the input images and the other with the highly unbalanced number of EMIC- and NOEMIC-type images. Such an unbalance was already expected to happen, since EMIC waves are a very spatially localized and sporadic phenomenon (Saikin et al. 2015; Wang et al. 2015), which is why a way larger number of NOEMIC-type images are present in our data set. Another important aspect of the BoF technique that favors its usage is that it is quite accurate when excluding the unwanted cases, which in, this context, are the images contained in the NOEMIC class. This characteristic helps us to significantly narrow down the number of analyzed images in larger data sets, saving a great deal of time. Therefore, we argue that the BoF technique is performing nearly as good as the visual classification method, which gives us confidence to apply it in other data sets.

6.2. Likely EMIC Wave Event Candidates According to the BoF Technique

The set of 36,014 unclassified images spanned the time interval from 2014 August 26 to 2016 December 31. From the

Visually classified as EMIC x BoF classified as NOEMIC



Visually classified as NOEMIC x BoF classified as EMIC




Figure 7. An example of wrong BoF classifications for both RBSP-A (panels (a) and (c)) and RBSP-B (panels (b) and (d)) data. Each green circle represent an extracted feature of the image as obtained by the SURF detector. Panels (a) and (b) show examples of images that have been visually classified as likely EMIC cases, therefore pertaining to the EMIC class, but the BoF classification defined them as belonging to the NOEMIC class. The opposite situation occurs in panels (c) and (d).

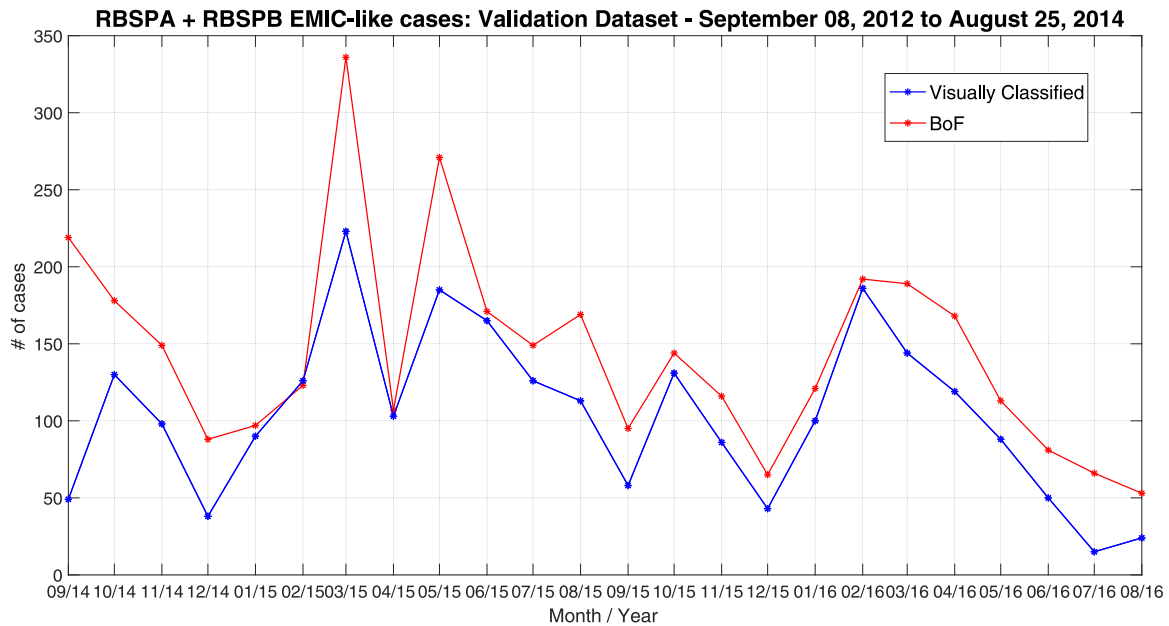


Figure 8. Comparison of BoF’s performance against that of the visual classification method. The validation data set is used and the period it refers to is shown on the top of both panels. The monthly number of EMIC cases as given by the visual classification (blue line) and by the BoF classification (red line) methods is shown in the vertical axis.

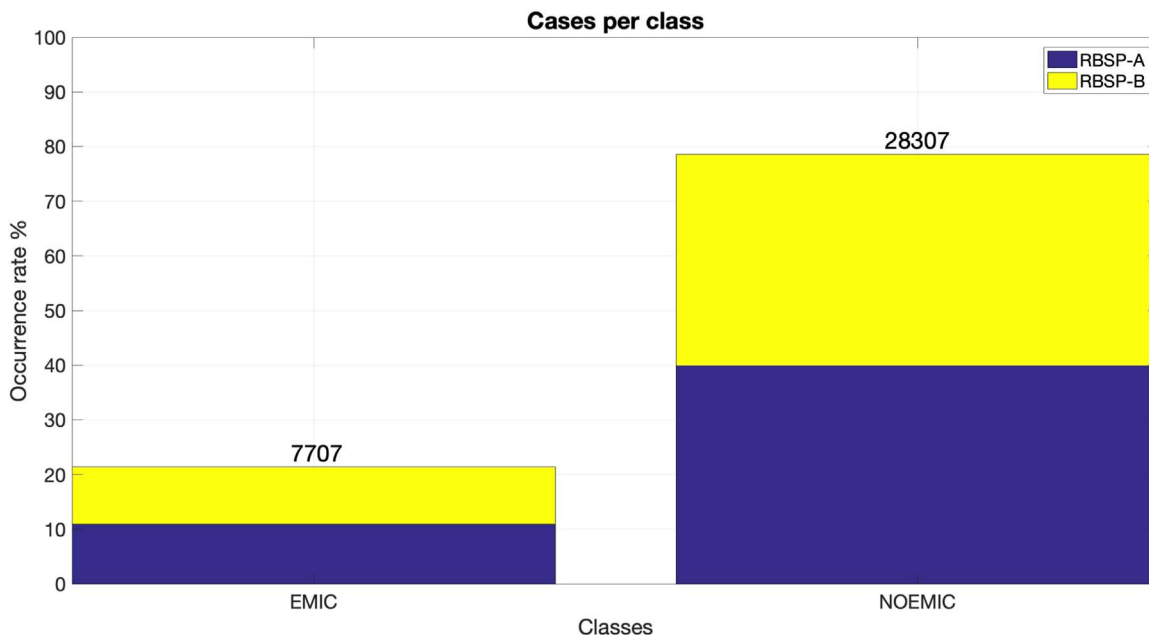


Figure 9. BoF’s classification of unclassified data set (36,014 images). About 21% of the input images are classified as EMIC and 79% are classified as NOEMIC. RBSP-A and RBSP-B observed almost the same number of cases per class.

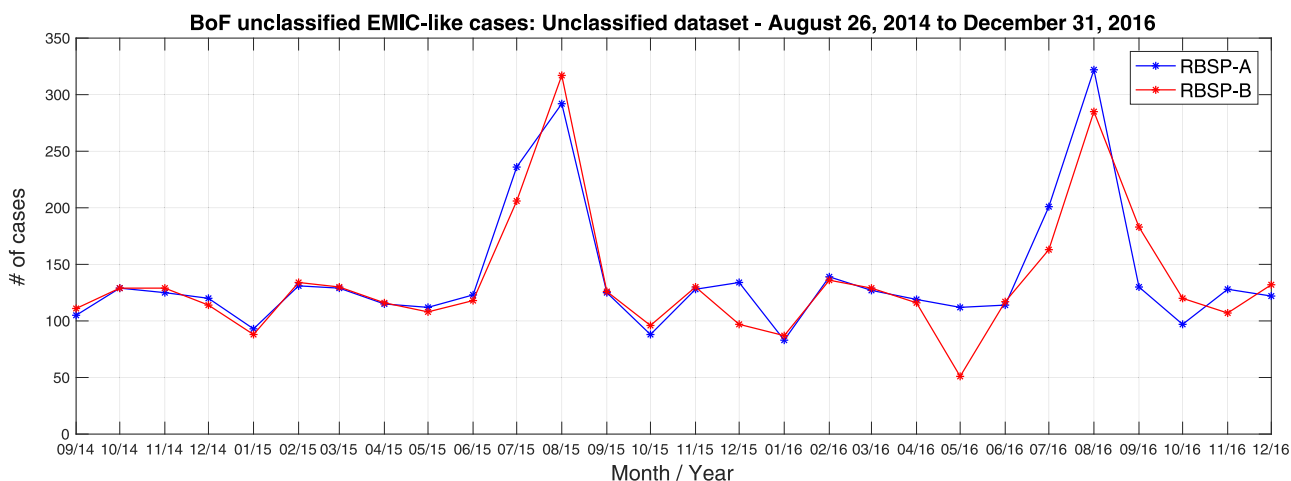


Figure 10. Monthly number of likely EMIC wave events as determined only by the BoF technique for the unclassified data set, which contains 36,014 stacked up spectrogram images. Data from RBSP-A (blue line) and RBSP-B (red line) is shown.

beginning of this period up to about 2016 March 30, the Van Allen Probes’ orbit apogee has covered all MLT sectors for the second time since the start of the mission in 2012 September. From early 2016 April to 2016 December 31, the satellites have covered the 11:00 → 06:00 → 00:00 → 19:00 MLT sectors, as will be discussed below. The BoF technique has been applied to this whole data set and the results are summarized both in Figures 9 and 10. The technique predicted that about 21% (7707) of the images should be categorized as EMIC cases, whereas the remaining 79% (that is, 28,307 images) should be classified as NOEMIC. Figure 9 also tells us that the number of either EMIC or NOEMIC cases is almost balanced between the two satellites, as also confirmed by Figure 10. The exact amount of either EMIC or NOEMIC cases that are either visually classified or classified by the BoF technique using in situ data from each Van Allen Probes is shown in Table 1 for the training, validation, and unclassified data sets. The number

of likely EMIC wave event candidates is nearly equal for RBSP-A and RBSP-B.

6.3. Spatial Distribution of EMIC Wave Events as Obtained by the BoF Technique

With exception of the training data set that contains 2000 images, we applied the BoF technique in the remaining data, i.e., 64,204 stacked up spectrogram images, to estimate the number of EMIC wave event cases during the entire span of our data set. The images grouped in the EMIC class by the BoF technique are split up according to the precessions of the Van Allen Probes’ orbits (Table 4), and their spatial distribution is shown in Figure 11 as a function of MLT and the spacecraft’s geocentric distance. Notice that the bin’s size in radial distance shrinks at farther out distances since they follow the same binning procedure in Figure 1. The four bins in radial distance are 3–4 R_E , 4–5 R_E , 5–5.5 R_E , and 5.5–6 R_E . In Figure 11, we do *not*

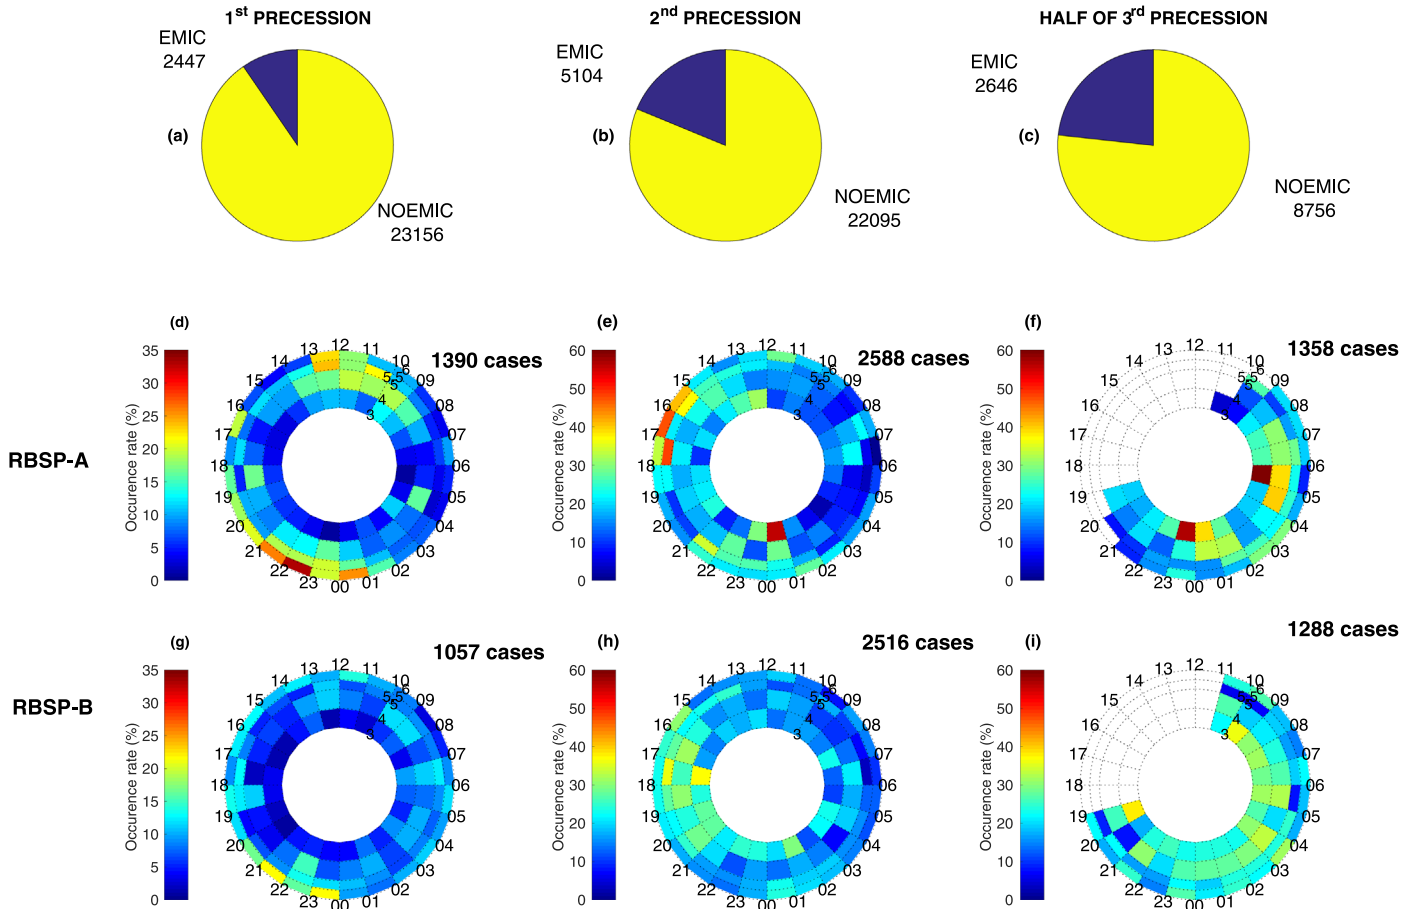


Figure 11. Distribution of images per classes according to the BoF classification process for (a) the first, (b) the second, and (c) the beginning of the third precession. In all intervals, the NOEMIC class contained most of the cases; see Table 4. Panels (d), (e), (f), (g), (h), and (i) contain the distribution of EMIC class cases for MLT for RBSP-A and RBSP-B, respectively.

Table 4

The Number of Images, per Van Allen Probe (Either A or B), Classified as EMIC or NOEMIC via the BoF Technique

Data Set	EMIC		NOEMIC		Total
	A	B	A	B	
First precession	1390	1057	11,461	11,695	25,603
Second precession	2588	2516	11,068	11,027	27,199
Half of the third precession	1358	1288	4516	4240	11,402
Total	5336	4861	27,045	26,962	64,204

Note. The data set used for the BoF classification is comprised of both validation and unclassified data sets, and they are divided into three precessions of the Van Allen Probes’ orbits. The first precession interval goes from 2012 September 8 to 2014 June 19, whereas the second and half of the third precessions go from 2014 June 20 to 2016 March 29 and from 2016 March 30 to 2016 December 31, respectively.

discriminate between likely EMIC events that are found either during the inbound or outbound portions of the orbit. The left, middle, and right columns of Figure 11 refer to the spatial distribution of likely EMIC wave events during the first, second, and half of the third precessions, respectively. The occurrence rate of likely EMIC wave events is calculated as follows. At a given MLT and radial distance bin, say 11:00–12:00 MLT and 3–4 R_E , we count the number of images N_I acquired at that bin during the analyzed interval, e.g., if the period refers to the first





precession, then the analyzed interval would be between 2012 September 8 to 2014 June 29. Then, for that same bin, we determine the number N_E of images classified by the BoF technique as likely EMIC wave events, thus the occurrence rate at that bin is given as N_E/N_I multiplied by 100 in order to provide a percentage value.

The spatial (BoF-based) EMIC wave distribution in the first precession, as presented in Figures 11(d) and (g), shows a slightly intensified occurrence rate near the noon (12:00 MLT) and midnight sectors, as has already been found in the literature (see, e.g., Saikin et al. 2016). The second precession contained most of the EMIC cases in our data set and both the dusk-side and part of the night-side sectors show the higher occurrence rates, as one can see in panels (e) and (h). Investigating whether the incidence of geomagnetic storms during the second precession was higher or not as compared to the previous precession could perhaps provide an explanation for an increased occurrence rate in these MLT sectors, since during geomagnetically perturbed periods night-side particle injections that contribute with storm-time ring current enhancements are expected to be higher, thus increasing the rate of EMIC wave generation in the midnight-dusk sector (see, e.g., Thorne 2010; Tetrick et al. 2017). This investigation, however, lies outside of the scope of the present work. During the period encompassing half of the third precession, almost all EMIC cases are concentrated between the midnight and dawn-side sectors, i.e., between 00:00 MLT and 06:00 MLT, as shown in panels (f) and (i). The relatively higher occurrence rate compared to the

preceding precessions is due mostly to the fact that a smaller number (roughly a factor of 2) of images have been analyzed in half of the third precession. An analysis of the interplanetary medium data should provide insights on the likely causes for the intensified EMIC wave occurrence rate on the dawn-side–midnight sector, as predicted by the BoF technique. We leave this step to a future study though.

The authors thank the EMFISIS’ instrument team for providing scientific data which is publicly available at the website <https://emfisis.physics.uiowa.edu/data/index>. V.M.S. acknowledges the financial support from the Brazilian National Council for Research and Development (CNPq) PCI grants 301228/2019-1 and 300053/2020-7.

ORCID iDs

Claudia Medeiros  <https://orcid.org/0000-0002-6914-5799>
 V. M. Souza  <https://orcid.org/0000-0001-7294-7963>
 L. A. Da Silva  <https://orcid.org/0000-0002-8822-5030>
 P. R. Jauer  <https://orcid.org/0000-0002-8064-5030>

References

- Alves, L. R., Da Silva, L. A., Souza, V. M., et al. 2016, *GeoRL*, **43**, 978
 Bay, H., Ess, A., Tuytelaars, T., & Gool, L. V. 2008, *Computer Vision and Image Understanding*, **110**, 346
 Camporeale, E., Wing, S., & Johnson, J. R. 2018, *Machine Learning Techniques for Space Weather* (Amsterdam: Elsevier)
 Clausen, L. B. N., Baker, J. B. H., Ruohoniemi, J. M., & Singer, H. J. 2011, *JGRA*, **116**, A10205
 Clilverd, M. A., Rodger, C. J., McCarthy, M., et al. 2017, *JGRA*, **122**, 534
 Cornwall, J. M. 1965, *JGR*, **70**, 61
 Dietterich, T. G., & Bakiri, G. 1995, *J. Artificial Intelligence Research*, **2**, 263
 Duda, R. O., & Hart, P. E. 1973, *Pattern Classification and Scene Analysis* (New York: Wiley)
 Halford, A. J., Fraser, B. J., Morley, S. K., Elkington, S. R., & Chan, A. A. 2016, *JGRA*, **121**, 6277
 Haykin, S. 2008, *Neural Networks and Learning Machines* (3rd ed.; New Jersey: Pearson Education)
 Horne, R. B., Lam, M. M., & Green, J. C. 2009, *GeoRL*, **36**, L19104
 Jaynes, A. N., & Usanova, M. E. 2020, *The Dynamic Loss of Earth’s Radiation Belts* (Amsterdam: Elsevier)
 Kennel, C. F., & Petschek, H. E. 1966, *JGR*, **71**, 1
 Kletzing, C. A., Kurth, W. S., Acuna, M., et al. 2013, *SSRv*, **179**, 127
 Li, X., Baker, D. N., Temerin, M., et al. 2005, *SpWea*, **3**, 04001
 Li, Z., Millan, R. M., Hudson, M. K., et al. 2014, *GeoRL*, **41**, 8722
 Mauk, B. H., Fox, N. J., Kanekal, S. G., et al. 2012, *SSRv*, **179**, 3
 Medeiros, C., Souza, V. M., Vieira, L. E. A., et al. 2019, *ApJ*, **872**, 36
 O’Hara, S., & Draper, B. A. 2011, arXiv:1101.3354
 Ozeke, L. G., Mann, I. R., Murphy, K. R., Sibeck, D. G., & Baker, D. N. 2017, *GeoRL*, **44**, 2624
 Reeves, G. D., McAdams, K. L., Friedel, R. H. W., & O’Brien, T. P. 2003, *GeoRL*, **30**, 1529
 Reeves, G. D., Morley, S. K., Friedel, R. H. W., et al. 2011, *JGRA*, **116**, 2213
 Remya, B., Tsurutani, B. T., Reddy, R. V., Lakhina, G. S., & Hajra, R. 2015, *JGRA*, **120**, 7536
 Roederer, J. G., & Zhang, H. 2016, *Dynamics of Magnetically Trapped Particles* (Berlin: Springer)
 Saikin, A. A., Zhang, J.-C., Allen, R. C., et al. 2015, *JGRA*, **120**, 7477
 Saikin, A. A., Zhang, J.-C., Smith, C. W., et al. 2016, *JGRA*, **121**, 4362
 Shprits, Y. Y., Chen, L., & Thorne, R. M. 2009, *JGRA*, **114**, A03219
 Souza, V. M., Lopez, R. E., Jauer, P. R., et al. 2017, *JGRA*, **122**, 10084
 Summers, D., & Thorne, R. M. 2003, *JGRA*, **108**, 1143
 Tetrick, S. S., Engebretson, M. J., Posch, J. L., et al. 2017, *JGRA*, **122**, 4064
 Thorne, R. M. 2010, *GeoRL*, **37**, L22107
 Thorne, R. M., & Horne, R. B. 1992, *GeoRL*, **19**, 417
 Thorne, R. M., Horne, R. B., Jordanova, V. K., Bortnik, J., & Glauert, S. 2013, *GMS*, **169**, 213
 Usanova, M. E., Drozdov, A., Orlova, K., et al. 2014, *GeoRL*, **41**, 1375
 Vapnik, V. N. 1999, *ITNN*, **10**, 988
 Wang, D., Yuan, Z., Yu, X., et al. 2015, *JGRA*, **120**, 4400
 Wing, S., Johnson, J. R., Camporeale, E., & Reeves, G. D. 2016, *JGRA*, **121**, 9378
 Zhang, J., Halford, A. J., Saikin, A. A., et al. 2016a, *JGRA*, **121**, 11086
 Zhang, J.-C., Kistler, L. M., Mouikis, C. G., et al. 2010, *JGRA*, **115**, 6212
 Zhang, X.-J., Li, W., Ma, Q., et al. 2016b, *JGRA*, **121**, 6620