Precise and Accurate Short-term Forecasting of Solar Energetic Particle Events with Multivariate Time Series Classifiers

Sumanth A. Rotti [1], Berkay Aydin [2] and Petrus C. Martens [1]

[1] *Georgia State University*
Department of Physics and Astronomy
Atlanta, GA, USA
[2] *Georgia State University*
Department of Computer Science
Atlanta, GA, USA

ABSTRACT

Solar energetic particle (SEP) events are one of the most crucial aspects of space weather that require continuous monitoring and forecasting using robust methods. We demonstrate a proof of concept of using a data-driven supervised classification framework on a multivariate time series data set covering solar cycles 22, 23, and 24. We implement ensemble modeling that merges the results from three proton channels (E≥10 MeV, 50 MeV, and 100 MeV) and the long band X-ray flux (1–8Å) channel from the Geostationary Operational Environmental Satellite missions. Our task is binary classification, such that the aim of the model is to distinguish strong SEP events from nonevents. Here, strong SEP events are those crossing the Space Weather Prediction Center's "S1" threshold of solar radiation storm and proton fluxes below that are weak SEP events. In addition, we consider periods of non-occurrence of SEPs following a flare with magnitudes ≥C6.0 to maintain a natural imbalance of sample distribution. In our data set, there are 244 strong SEP events comprising the positive class. There are 189 weak events and 2,460 "SEP-quiet" periods for the negative class. We experiment with summary statistic classifier, one-nearest neighbor and supervised time series forest (STSF), and compare their performances to validate our methods for prediction windows from 5 min up to 60 min. We find STSF to perform better under all circumstances. For an optimal classification threshold of ≈0.3 and a 60 min prediction window, we obtain: TSS = 0.850, HSS = 0.878, GSS = 0.783.

Corresponding author: Sumanth A. Rotti
srotti@gsu.edu

1. INTRODUCTION

The activity of the Sun is considered the prime source of space weather (SWx) that constitutes different eruptive phenomena including solar flares (SFs) and coronal mass ejections (CMEs). SF is a sudden brightening in the solar atmosphere in a coronal soft X-ray source that is observed to have close relationships with CMEs (Feynman & Hundhausen 1994) and other observable phenomena such as jets and filament eruptions (Schrijver & Siscoe 2010). CMEs are clouds of plasma formed in the lower corona that often move faster (>1000 km.s^{-1}) than the ambient solar wind (Low 1996; Chen 2011). Another manifestation of solar activity constituting the emission of energetic electrons, protons, and heavier ions from the Sun is called a solar energetic particle (SEP) event (Klein & Dalla 2017). Typically, large flares and shock fronts of fast CMEs are known to accelerate SEPs and are often considered precursors or parent eruptions (Cane et al. 1986; Kahler 1992; Gopalswamy et al. 2001; Marqué et al. 2006; Gopalswamy et al. 2008; Swalwell et al. 2017; Gopalswamy et al. 2017; Cliver & D'Huys 2018; Rotti & Martens 2023).

The Sun releases ions of about tens of Mega electron-Volt (MeV) and more in many SEP events. The motion of such highly charged particles is dictated by magnetic field lines (Reames 2013). To be geo-effective, SEPs should reach the near-Earth space on a magnetically well-connected path (Cane et al. 1988). Generally, it is understood that the eruptions at the western side of the Sun have a higher probability of SEPs being geo-effective due to the spiral structure of the interplanetary magnetic field lines, known as the Parker spiral (Parker 1965; Reames 1999). Such Earth-bound SEPs are hazardous on many levels. The impacts include severe technological (Smart & Shea 1992) and biological effects on various economic scales (Schrijver & Siscoe 2010). Although Earth's magnetic field provides us a protective shield from the energetic particles and filters them out from reaching the ground, they can be life threatening for humans on missions outside of the Earth's magnetosphere and aircraft travels along polar routes (Beck et al. 2005; Schwadron et al. 2010). For instance, long-lasting strong SEP events pose a radiation hazard to astronauts and electronic equipment in space (Jiggens et al. 2019). In addition, the particles getting diverted to Earth's magnetic poles disturb the ionosphere's upper layers, causing disruption to high-frequency radio, GPS, and satellite communications (Desai & Giacalone 2016).

Over the last four decades, great progress in space exploration has provided near-continuous observations of solar activity from a fleet of advanced space-based satellites (Usoskin 2017). These observational data should be analyzed in operational contexts to mitigate SWx effects on our human explorers and technological systems (Jackman & McPeters 1987). Therefore, we require robust tools and systems to forecast eruptive event occurrences such as SEPs and send warning signals before the event. Researchers across the globe have been focusing on implementing a variety of model-driven techniques for predicting SEP events, mostly concentrating on predicting the peak flux-related characteristics. To forecast SEP event occurrences, physics-based

and data-driven statistical models have been designed based on the parameters of parent solar eruptions such as SFs and CMEs (Van Hollebeke et al. 1975; Posner 2007; Kahler et al. 2007; Balch 2008; Laurenza et al. 2009; Núñez 2011; Falconer et al. 2011; Dierckxsens et al. 2015; Anastasiadis et al. 2017; Alberti et al. 2017; Papaioannou et al. 2018).

Machine learning (ML) methods have been at the forefront of SEP event forecasting in the last decade due to their success in many other areas of research and operations (Boubrahimi et al. 2017; Swalwell et al. 2017; Engell et al. 2017; Camporeale 2019; Aminalragia-Giamini et al. 2021; Lavasa et al. 2021; Whitman et al. 2022; Hosseinzadeh et al. 2024). Our previous work (Rotti, Aydin & Martens 2024) implemented time series classifiers on a multivariate time series (MVTS) constituting solar X-ray and proton fluxes data set of SEP events and showed that the interval-based supervised time series forest (STSF) model is more efficient in the classification of strong SEP events. Possible areas of improvement here are to consider an actual class imbalance scenario, that is, including the phases of non-SEPs and comparing with baseline time series classification algorithms for the task at hand. Expanding on these ideas in the present work, we implement STSF, summary statistic classifier (SSC) and one-nearest neighbor (1NN) classifiers on an extended data set constituting SEP-quiet periods and analyze their performances. With that, we compare the results and establish a proof of concept of the model to be implemented in a near-real-time (NRT) scenario. That is, we aim to build a high fidelity (low-risk) and interpretable short-term (low lead time) predictive model for reliable SEP event forecasting systems suitable to operational standards.

SEPs observed in near-Earth space are rare events compared to major flares and CMEs. In addition, the occurrences of SEPs are dependent on precursor solar eruptions (Cliver et al. 2022; Rotti & Martens 2023). Thus, it makes physical sense to build SEP predictive models based on the parameters of parent solar events. However, in the present work, we consider only SFs and defer to the use of CME data due to their low data quality, such as the projection effects and difficulty of tracking them as 3d objects.

The rest of the paper is organized as follows: Section 2 provides information about our data set and data preparation steps used in this work. Section 3 presents our research methodology, including the descriptions of the time series classification models and data partitioning scheme. Section 4 discusses the training and validation phases of the model and presents the experimental evaluation framework. Lastly, Section 5 provides a summary of our work and future endeavors.

2. DATA

In this work, we implement MVTS data integration, model training and selection pipeline toward building an efficient architecture for short-term SEP event forecasting. We consider a binary classification task between strong SEP and weak or non-SEP

events. In the present framework, we consider two classes, "SEP" and "SEP-quiet," where SEP-quiet periods include smaller proton flux increases, as well as the flares that did not lead to SEP events near Earth. With that, we maintain a natural class imbalance between the occurrence and non-occurrence of strong SEP events. Our methodology for the creation of SEP-quiet samples will be presented in the next Section. The strong and weak SEP samples in our work come from the Geostationary Solar Energetic Particle (GSEP) events data set (Rotti et al. 2022), which is discussed later in Section 2.4.

2.1. *SEP-quiet periods*

The National Oceanic and Atmospheric Administration (NOAA) has been operating the Geostationary Operational Environmental Satellite (GOES) series in geosynchronous orbits since 1976 (Sauer 1989; Bornmann et al. 1996). The space environment monitor (SEM) onboard GOES missions measures solar radiation in the X-ray and EUV region and the in-situ magnetic field and energetic particle environment (Grubb 1975). In this work, we utilize the GOES solar X-rays (Garcia 1994) and energetic proton fluxes (Onsager et al. 1996). We create our negative samples using "SEP-quiet" periods in two stages: (i) obtaining lists of SF and SEP events and (ii) gathering the corresponding GOES flux measurements for each event. We note that we take into account the scaling factors initially implemented by NOAA in the GOES-08 through -15 X-ray measurements. Towards that, we updated our data set by removing the necessary scaling factors as per the NOAA guidelines[1] to get true X-ray fluxes.

NOAA-GOES has traced flaring activity since the beginning of 1974 and offers a catalog[2] with spatial and temporal specifications, flare magnitude, and associated active region (AR) information. However, some flare locations and/or AR numbers are missing from the GOES catalog or probably have known errors (Milligan & Ireland 2018; Angryk et al. 2020; Rotti et al. 2020). NOAA categorizes SFs based on soft X-ray peak flux in the wavelength range 1–8 Å. Flare classes from least to most intense are labeled as A, B, C, M, and X, where, each category indicates an increase in flare intensity on a logarithmic scale.

2.2. *Non-SEP samples*

We select our non-SEP samples such that there are no enhancements or relatively no enhancements in the GOES proton fluxes following flaring episodes. In other words, we identify large flares that do not lead to significant variations in the GOES ≥ 10MeV proton fluxes relative to the background levels. This step helps us increase the sample size over the previous study. At first, we chose M and X-class flares from the NOAA flare list between 1986 and 2018 that have not been associated with strong

---

[1] GOES XRS scaling factors: https://ngdc.noaa.gov/stp/satellite/goes/doc/GOES_XRS_readme.pdf
[2] The NOAA flare list: https://www.ngdc.noaa.gov/stp/space-weather/solar-data/solar-features/solar-flares/x-rays/

Figure 1. Distribution of soft X-ray flare peak intensities based on GOES flare classification for SEP-associated flares in the GSEP data set. There are a total of 383 SEP with flare class information, of which 59 events have flare magnitudes less than GOES-class C6.0.

SEP events. Then, to add a few more negative samples, we extend our lower limit of a "non-SEP" flaring intensity using a data-driven approach. The GSEP data set consists of carefully identified SF information for 383 associated SEP events. In that, we see a distribution of strengths amongst 59 SEP events associated with <C6.0 flares, where 38 are weak events and 21 are strong events. Nonetheless, we inspect all of the SEP-associated flares as shown in Figure 1 and find that <C6.0 flares correspond to the bottom 15th percentile. Hence, we consider C6.0 as our bottom threshold for flares. Then, we add all flares from M1.0 to C6.0 to our earlier flare list and obtain a total of 7,981 flares. We filter the preliminary negative-sample flares in two steps: (1) we do not consider flares during an ongoing SEP event (either rise or decay) and (2) we remove all consecutive flares within 11 hr following the onset of the flare at consideration. With this filtering step, our final count of non-SEP flares with an intensity ≥C6.0 reduced to 2,460. In Figure 2, we show the distribution of the number of "SEP-quiet" flares in each sub-divisions of GOES (2a) C-, (2b) M-, and (2c) X-class.
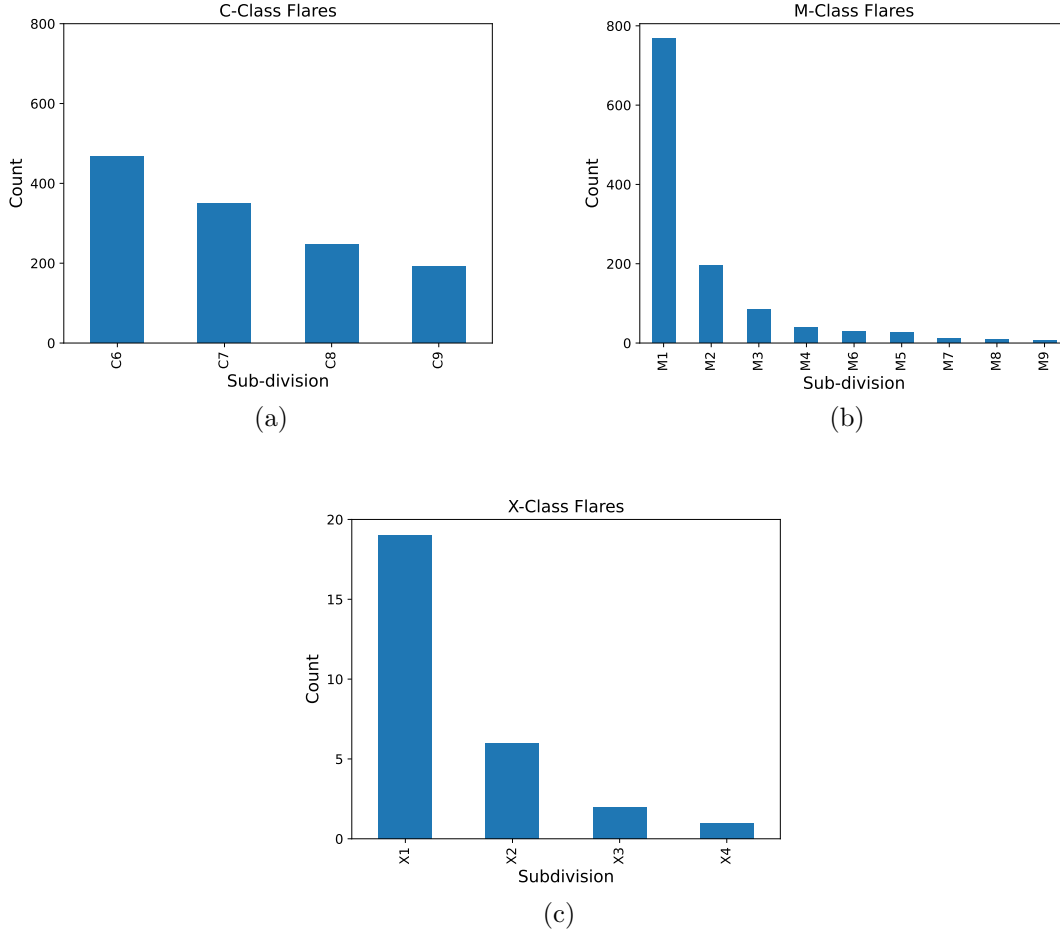
Whitman et al. (2022) recently organized an SEP model validation (SEPVAL) challenge in 2023[3]. For this, two lists of 33 SEP and 30 non-SEP events from solar cycles (SCs) 24 and 25 have been provided. In this regard, we cross-verified our non-SEPs with the SEPVAL's list of 15 non-SEP events for cycle 24. We do not include any (SEP or non-SEP) events from SC25 in the present work to make our analysis comparable with earlier studies.

2.3. Time Series Slicing

To generate our dataset, we have interpolated and integrated particle and X-ray flux data from several GOES missions. First, we use the one-minute averaged X-ray (1–8 Å) fluxes measured by the X-ray sensor (XRS) onboard GOES. The archived data is available online from NOAA's website[4]. In addition to X-ray data, we use the following five-minute averaged integral proton fluxes measured by the SEM suite: (1) E≥10 MeV fluxes corresponding to P3, (2) E≥50 MeV fluxes corresponding to P5,

---

[3] SEPVAL Challenge: https://ccmc.gsfc.nasa.gov/community-workshops/ccmc-sepval-2023/
[4] GOES-XRS: https://www.ncei.noaa.gov/data/goes-space-environment-monitor/access/avg/

**Figure 2.** Distribution of the GOES subdivisions of (a) C- (b) M- and (c) X-class flares present on our "SEP-quiet" periods.

and (3) E≥100 MeV fluxes corresponding to P7. We interpolate the proton fluxes to one-minute averages and retain the time stamp of every first measurement. The rationale here is to capture the flare dynamics that occur over a few minutes. In the present work, we follow the same procedure as discussed in Rotti et al. (2024) to generate our event slices for the additional negative samples and ensures that our analysis and models' predictive capability are relevant and comparable across different SEP event prediction tasks.

We consider flares as the precursor to the onset of an SEP event. In our earlier statistical study on strong and weak SEP events (Rotti & Martens 2023), we find that 77(±4)% of SEP event onsets occur within 11 hrs after the associated flaring eruption. However, sampling flares from the onset will likely lead the classifier to memorize the pattern of finding a spike at the start of the MVTS instance. A simple solution to randomize the sampling is to include 1 hour of data prior to the flare onset because two to three flares may occur within an hour of observation in many instances. With that consideration, the classifiers become free from memorizing to

find a spike at the start of the data or at the end of 60 data points corresponding to one hr. We accommodate both these aspects while creating our non-SEP time series sampling because our goal is to develop a highly accurate model predicting strong SEP events. As reported in our previous work (Rotti et al., 2024), we note that our modeling schema can predict well-connected events in addition to strong proton enhancements. A caveat/drawback of the model would be to classify events that fluctuate around the 10 pfu threshold.

Based on our filtered list of 2,460 flares for negative samples, we consider 1 hour of data (all four fluxes mentioned above) before the flare onset and 11 hr afterward for each flaring event (as shown in Figure 3). Therefore, the total length of each MVTS corresponds to 12 hr. As our final data set comes with a cadence of one minute, each MVTS is a size-720 sequence of soft X-ray and integral proton fluxes. All time series data was cleaned for data gaps with linear interpolation and spike removal by replacing what would be considered an outlier (spike) with an interpolated value from moving averages. We note here that the fluxes are predominantly derived from primary GOES missions. Nonetheless, when there are ambiguous measurements in the primary satellite, we consider data from the secondary satellite of the corresponding period of observation.

Because we are focusing on short-term predictions, we explore multiple prediction windows: 5-, 15-, 30-, 45-, and 60-min. In each case, we exclude the respective minutes of observations from the 720-size MVTS at the end by maintaining a constant 11 hr observation window as input to the model. Such a continuous and consistent observation window is needed for training a predictive model to be established in a near-real-time operational environment. The prediction window indicates the corresponding time in the future at the end of the observation window for when the model output is valid.
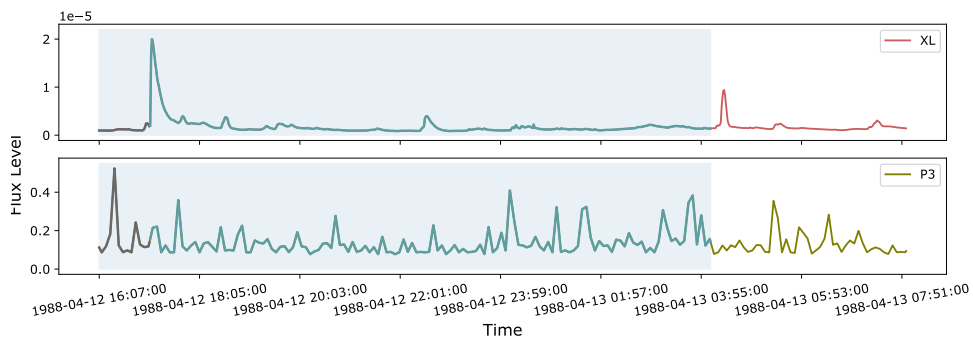
Figure 3. An example of the sampling technique for the non-SEP samples in our data set. We show the XL and P3 fluxes for representation only. The blue-shaded region in the plot shows the 12-hour length of the input time series of our sample. Here, the grey-shaded line on the left corresponds to the one hour of fluxes prior to the ($\geq$C6.0) flare onset, while the rest of the time series within the blue region covers 11 hours after the onset.

2.4. *GSEP data set*

Our primary source of SEP events comes from the recently published open-source GSEP data set (Rotti et al. Rotti et al. (2022)). It consists of a carefully vetted SEP events catalog with associated parent eruption metadata and time series slices. The data set comprises 433 (- 244 strong and - 189 weak) SEP events covering solar cycles 22 to 24. According to the NOAA Space Weather Prediction Center (SWPC), a strong SEP event is defined as having proton intensities $I_P \geq$ 10 pfu (1 pfu = 1 particle per cm^2.s.sr) in the GOES five minutes averaged E$\geq$10 MeV integral energy channel for at least three consecutive readings (Bain et al. 2021). In the GSEP data set, "weak" SEP events correspond to an enhancement in proton intensities in E$\geq$ 10 MeV channel whose peak flux measurements do not reach to 10 pfu. The description of the data set and its development can be found in Rotti et al. (2022) and Rotti & Martens (2023).

3. METHODOLOGY

We consider the term 'SEP events' analogous to solar protons events (SPEs). While variations exist, positive labels are usually associated with the occurrence of strong SEPs based on the integral proton fluxes recorded by the GOES-P3 channel crossing the 10 pfu threshold. As mentioned earlier, we have 189 weak SEP events in the GSEP data set and we add 2,460 non-SEPs to increase the negative samples to 2,649. For the positive class, we have 244 strong SEP events. The total size of the data set is 2,893 samples with a class-imbalance ratio of $\approx$ 1:11. Figure 4 shows the total number of samples in our data set from 1986 to 2018. Here, 'Events" are defined as those SEPs crossing the SWPC threshold of 10 pfu in the E$\geq$10 MeV channel, and "Non-Events" constitute both weaker enhancements of <10 pfu and no enhancements. In Section 3.1, we briefly explain the method of time series classification implemented in the present work.

3.1. *Time Series Classification*

We examine the performance of time series classifiers in the framework of a binary classification task for predicting SEP events. Time series classification uses supervised ML to analyze labeled classes of time series data and then predicts the class to which a new instance belongs. This is important in SWx predictions, where observational data is analyzed to support NRT operations and classification accuracy becomes critical. Furthermore, for short-term predictions, the algorithm must be highly accurate and robust to be useful. In our previous work (Rotti et al. 2024), we showed that an ensemble framework of interval-based classifiers achieved high performance in classifying between strong and weak SEP events.

In the present work, we implement three time series classifiers and compare their performance to establish certainty in our SEP event prediction system. The motive is to see whether the input time series span leads to a strong SEP event (positive class) or not (negative class). We implement the same column ensemble schema presented

Figure 4. The distribution of SEP and non-SEP samples in our data set between 1986 and 2018. We split our data set into three parts of non-overlapping years: Partition 1 = 1986 - 1992; Partition 2 = 1993 to 2002; and Partition 3 = 2003 to 2018. Here, each split consists of 998, 974 and 921 samples, respectively. The data set has a cumulative class-imbalance ratio of ≈ 1:11. In the legend, "Events" are defined as those SEPs crossing the SWPC threshold of 10 pfu in the E≥10 MeV channel, and "Non-Events" consists of both weaker enhancements <10 pfu or no enhancements.

in our previous work (Rotti et al. 2024), where the ensemble estimator allows multiple feature columns of the input to be transformed separately. The statistical features generated by each classifier on samples of the original time series are ensembled to create a single output. Each feature is assigned a score that indicates how informative it is towards predicting the target variable (Hansen & Salamon 1990; Schapire 1990; Arbib 2003). The predictions from individual classifiers are then aggregated with equal votes using prediction probabilities. We will give a brief overview of our models in the following sections.

3.2. *Summary Statistic Classifier*

The feature-based summary statistic classifier is the first of the two baseline models we have used in this work. This classifier applies a time series feature transformation to an entire univariate time series using simple summary statistics and builds a random forest using these summary statistics. The summary statistics extracted are mean, standard deviation, minimum, maximum, and quantile values (at 0.25, 0.5, and 0.75, corresponding to first and third quartiles and median). The tabulated summary statistics are then fed to a random forest classifier for training a univariate time series classifier. For applying this method to our multivariate time series classification task, we employed a late fusion approach, where univariate time series classification models trained on individual parameters are ensembled based on a majority voting schema.

3.3. *One Nearest Neighbor*

The similarity-based (nearest neighbor) time series classification is one of the most popular classification methods due to its simplicity, nonparametric nature, potential for interpretability, and flexibility in utilizing different similarity schemes/metrics. It assigns a label for a new instance based on the target value of k-most similar instances in the training dataset. The similarity is assessed based on the inverse of a designated time series distance measure and utilization of elastic distance measures

allows capturing the local and global similarity patterns (Faouzi 2022). In this work, we implement a k-nearest neighbor (kNN, with k=1) classifier with Dynamic Time Warping (DTW) distance, which is a commonly used elastic measure (Sakoe & Chiba 1978). This 1NN is a baseline classifier for time series classification and similar to our earlier approaches, we use 1NN with DTW distance to issue predictions.

3.4. *Supervised Time Series Forest Classifier*

The supervised time series forest (STSF; Cabello et al. 2020) classifier comes under the category of interval-based models, which uses three representations (time, frequency, and derivative) of the input time series to extract features from a supervised selection of intervals. That is, for each representation, STSF builds an ensemble of decision trees on intervals selected through a supervised process wherein the algorithm finds the discriminatory intervals. STSF computes the region of interest (ROI) to highlight the location of discriminatory intervals, which are the intersected regions of such intervals. It extracts seven features, namely mean (μ), standard deviation (σ), slope (m), median, interquartile range (IQR), minimum value and maximum value, from each interval. The ranking of the interval feature is obtained by a scoring function that indicates how well the feature separates one class of time series from the other classes. The final set of intervals is obtained in a top-down approach to represent the entire series. The feature set is concatenated to form a new data set upon which decision trees are built. The final output is based on majority voting of averaged probability estimates of the individual estimators in the ensemble.

3.5. *Data Partitions*

As we are implementing a classification task with target labels, we split our data set into three sets: the training set, the holdout or validation set, and the test set. This partitioning schema is relevant to estimate the predictive capacity of the model on different aspects. For instance, the validation set is used in the present work to obtain the optimal threshold of classification for each model, estimate the skills scores and compare their performance. Furthermore, we use the unseen test set to sanity-check the best model's predictive capabilities.

We split our data set into three sets on non-overlapping years (1986-1992, 1993-2002 and 2003-2018) such that there are 998 samples for training, 974 samples for validation and 921 samples for testing the model. Here, each partition consists of a similar number of positive samples (≈ 80) and has similar class imbalance ratios, that is, $\approx 1:11$. Table 1 shows the number of data samples in each partition in relation to the binary class labels.

4. RESULTS

In this work, we aim to demonstrate the robustness of the methodology for SEP event forecasting that was proposed in our previous work (Rotti et al. 2024) and provide a proof of concept on its efficiency in generating short-term predictions. For

Table 1. Partitioning strategy of our data set.

	Training	Validation	Test
Positive	80	80	84
Negative	918	894	837

NOTE—Number of samples in each partition corresponding to the binary target labels. Here, labels correspond to a yes (positive) or no (negative).

this purpose, we compare three classifiers (SSC, 1NN and STSF) and analyze their performances. We construct a binary classification task such that the 'positive' class consists of strong SEP events while the 'negative' class contains the rest. We use the training set to train the model and perform a grid search for hyperparameter optimization of our classifiers. The best hyperparameters for STSF and SSC were on the default model settings with the number of estimators = 200.

We explore the model capabilities for different short-term prediction windows: 5-, 15-, 30-, 45-, and 60-minutes. We use the following terms henceforth to identify different prediction windows that act as lead times for model predictions, respectively: T_5, T_{15}, T_{30}, T_{45} and T_{60}. In each experiment, we consider the respective prediction window consisting of a constant 11-hour observation window from our MVTS samples before training the model. We implement multiple metrics (see Section 4.1) to estimate the predictive/forecast skill of our models.

4.1. Metrics

We consider a 2 × 2 contingency table with the following elements: true positive (TP), true negative (TN), false positive (FP), and false negative (FN). Here, TP indicates the number of correctly predicted strong SEP events (positive class), while TN represents the number of correctly predicted weak and non-SEP events (negative class). FP corresponds to the number of negative classes predicted as positives (false alarms), while FN corresponds to the number of positive class labels predicted as negatives (misses).

On a simple scale, statistical metrics such as precision (Eq. 1) and recall (Eq. 2) have been traditionally used to assess classifier performances. Precision for the positive class is used to evaluate the model's correct prediction with respect to the false alarms. Recall for the positive class characterizes the ability of the classifier to find all of the positive cases.

$$Precision = \frac{(TP)}{(TP + FP)} \tag{1}$$

$$Recall = \frac{(TP)}{(TP + FN)} \tag{2}$$

Focusing on the importance of positive classes, we consider the F_1-score for the positive class that can be estimated as the harmonic mean of precision and recall as shown in Equation 3. It ranges between 0 and 1 such that scores closer to 1 indicate the model to be better.

$$F_1 = 2 \times \frac{(Precision \times Recall)}{(Precision + Recall)} \qquad (3)$$

To understand the overall model performance, we use a 'weighted' average for the F_1-score (shown in Equations 4 and 5), which computes the score for each target class and uses sample weights that depend on the number of instances in that class while averaging. Here, i is the number of target classes in the data set, which is two in the present work.

$$F_{1_{weighted}} = \sum_{i=1}^{N} w_i \times F_{1_i} \qquad (4)$$

$$w_i = \frac{\text{Number of samples in class } i}{\text{Total number of samples}} \qquad (5)$$

We implement true skill statistics (TSS; Woodcock 1976; Dann 1985) to account for the false positive rate comparing the difference between the probability of detection and the probability of false detection as shown in Equation 6. TSS ranges from -1 to +1, where the latter indicates a perfect score. TSS ≤ 0 indicates it is worse than a random classification.

$$TSS = \frac{(TP \times TN) - (FP \times FN)}{(TP + FN) \times (FP + TN)} \qquad (6)$$

Furthermore, we consider the Heidke skill score (HSS; Heidke 1926) that measures the improvement of the forecast over a random prediction as defined in Equation 7. HSS with 1 indicates perfect performance and 0 indicates no skill. A no-skill means the forecast is not better than a random binary forecast based on class distributions.

$$HSS = \frac{2 \times ((TP \times TN) - (FP \times FN))}{((TP + FN) \times (TN + FN)) + ((FP + TN) \times (FP + TP))} \qquad (7)$$

The Gilbert Skill Score (GSS; Schaefer 1990) considers the number of hits due to chance, which is the frequency of an event multiplied by the total number of forecast events. This score formula is given by Equation 8. GSS ranges from -1/3 to 1. Here, 0 indicates no skill, while 1 is a perfect forecast.

$$GSS = \frac{TP - (\frac{(TP+FN) \times (TP+FP)}{TP+FP+TN+FN})}{(TP + FP + FN) - (\frac{(TP+FN) \times (TP+FP)}{TP+FP+TN+FN})} \qquad (8)$$

Accounting for TN to assess the performance of a binary class problem is essential in our context. Hence, we also implement Matthew's correlation coefficient (MCC)

as defined in Equation 9. MCC ranges from -1 to 1 where 0 indicates no skill and 1 shows perfect agreement between predicted and actual values. MCC is a robust metric that follows a comprehensive and balanced strategy for using the contingency table elements. That is, a high MCC score is obtained only when the predictions proportionately account for the size of both positive and negative samples in the data set. This is important because popular scores such as F_1 fall short to account for class imbalance.

$$MCC = \frac{(TP \times TN) - (FP \times FN)}{\sqrt{(TP + FP) \times (TP + FN) \times (TP + FP) \times (TN + FN)}} \qquad (9)$$
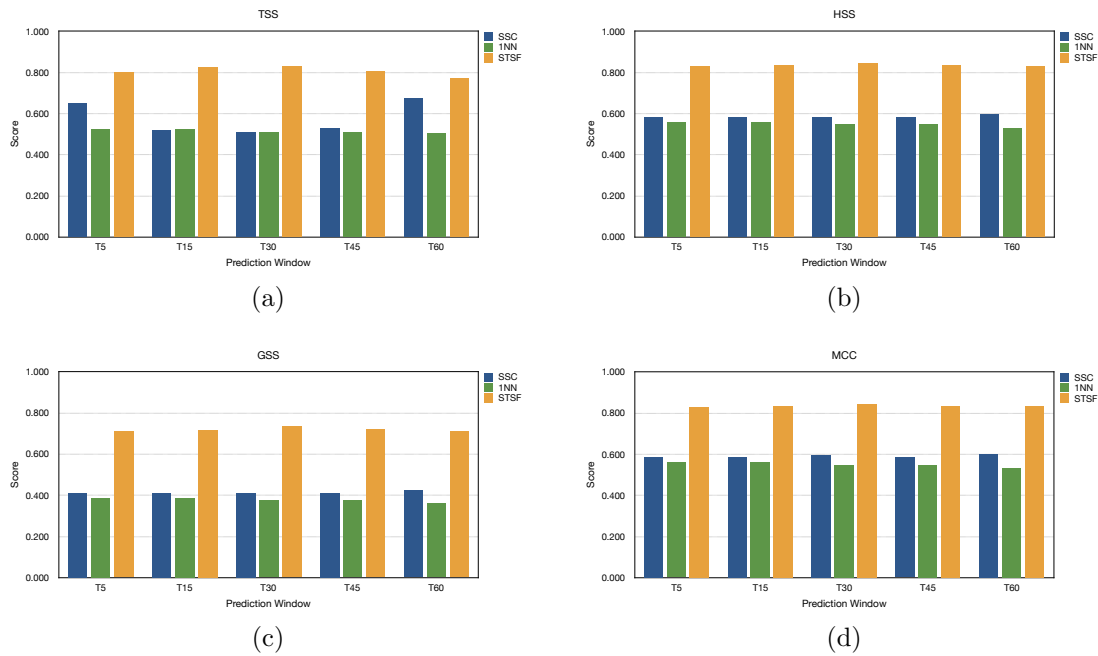
4.2. Validation set

As mentioned in Section 3.5, we use the second partition from our data set as a hold-out portion to estimate the optimal threshold for our classifiers. The classification threshold is the decision threshold that allows us to map the probabilistic output of a binary classifier to a binary category. In other words, it is a cut-off point used to assign a specific predicted class label for each sample. By default, the classification threshold in our models is 0.5. What that means is any prediction above 0.5 belongs to the positive class and that below 0.5 belongs to the negative class. However, 0.5 is not always optimal, especially when we have an imbalanced data set. Therefore, we identify a reliable threshold for the classifier that better discriminates between the two class labels. The first step is to derive the prediction probabilities from the validation set and extract the best trade-off between true positive rate (TPR) and false positive rate (FPR). In general, any binary classification model predicts mean probabilities for each input sample belonging to the positive class, where the prediction score from the classifier is greater than a parametrized threshold. Then, a classification threshold (from 0.0 to 1.0) is used to assign a binary label to the predicted probabilities.

$$J = Sensitivity + Specificity - 1 \qquad (10)$$

To find the optimal threshold that minimizes the difference between TPR and FPR, the Youden Index (J; Youden 1950) is used in the literature as defined in Equation 10. Here, sensitivity is the recall for the positive class and specificity is the recall for the negative class. J is the specialized version of TSS for the binary classification task. As a characteristic example, we show in Figure 5 the effect of "thresholding" on the performance of STSF by visualizing the variations in the skills due to changing thresholds. Here, the estimated optimal threshold of STSF for T_{30} on our data generation and sampling method is 0.28. We obtain optimal classification thresholds at all prediction windows for all the models and estimate the model skills. A comparison of the model performances is shown in Figure 6.

Computationally, the run time of SSC was the fastest, followed by STSF and then 1NN. In terms of performance, 1NN falls behind both models but maintains relatively close values for all the skill scores in each prediction window with $\sigma \approx 0.04$. On the other hand, SSC shows a similar trend in its skills with $\sigma \approx 0.01$ except for TSS ($\sigma \approx$

Figure 5. The variation in skills such as TSS, HSS, GSS and MCC with respect to increasing the classification threshold for the STSF model on the validation/hold-out set at a prediction window of thirty minutes (T_{30}). The optimal threshold value for the model is inset in the plot.

Figure 6. Comparison of SSC, STSF and 1NN classifiers with respect to variation in skills such as (a) TSS, (b) HSS, (c) GSS and (d) MCC for prediction windows T_5, T_{15}, T_{30}, T_{45}, and T_{60} on the validation set.

0.07). Although SSC was a little faster, the performance marginally lags on our data set compared to STSF. For example, the average TSS for SSC is $\approx 14\%$ lower than that for STSF. The STSF model performs exceptionally well compared to the other two baseline classifiers at all parameters. Specifically, the high scores at T_{60} (e.g., TSS = 0.850) provide satisfactory reliability to our methodology. Table 2 presents the contingency tables for all prediction windows corresponding to the optimal threshold of STSF on the validation set. The contingency tables for SSC and 1NN are provided in Tables 4 and 5, respectively, in the Appendix.

Table 2. Contingency tables for the STSF model on the validation set.

		$\mathbf{T}_5$ Predicted		$\mathbf{T}_{15}$ Predicted		$\mathbf{T}_{30}$ Predicted		$\mathbf{T}_{45}$ Predicted		$\mathbf{T}_{60}$ Predicted	
		Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak
True	Strong	65	15	67	13	67	13	65	15	62	18
	Weak	9	885	11	883	9	885	8	886	5	889

NOTE—Truth tables for the STSF model for different prediction windows. The first column is a shared entry of true labels against predictive labels for each experiment. The elements indicate the number of predictions with respect to the actual occurrences in the validation set.
prediction windows are shown by a subscript to T. For example, T_5 indicates a prediction window of 5 min.
Model name: STSF - Supervised Time Series Forest

Rotti et al.

Table 3. Contingency tables for the STSF model on the test set.

| | | $\mathbf{T}_5$ Predicted | | $\mathbf{T}_{15}$ Predicted | | $\mathbf{T}_{30}$ Predicted | | $\mathbf{T}_{45}$ Predicted | | $\mathbf{T}_{60}$ Predicted | |
		Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak
True	Strong	71	13	73	11	75	9	73	11	72	12
	Weak	6	831	6	831	8	829	6	831	6	831

Note.—Truth tables for the STSF model for different prediction windows. The first column is a shared entry of true labels against predictive labels for each experiment. The elements indicate the number of predictions with respect to the actual occurrences in the test set. Prediction windows are shown by a subscript to T. For example, T_5 indicates a prediction window of 5 min.
Model name: STSF - Supervised Time Series Forest

4.3. *Test set*

In our previous work (Rotti et al. 2024), we found that an ensemble of interval-based STSF classifiers is a highly efficient and optimal model to predict strong SEP events. STSF's approach to computing interval features in a supervised manner from three different time series representations by using robust statistics was a viable option for our problem specification. In this paper, we extended our modeling strategy by including non-SEP periods in our data set. Furthermore, we compared the performance of STSF and two baseline classifiers, namely SSC and 1NN. Based on the forecasting skill scores, we used the validation set from our data partitions to estimate the model performances. This resulted in STSF as the best model under all comparative paramters. We use the test set to further assess the model's capabilities.

In Table 3, we show the contingency or the confusion matrix elements derived for the STSF classifier on the test set for all prediction windows (T_5 to T_{60}) considered in this work. The increasing prediction window shows a relatively consistent type II error (false negative rate) except for T_{30}. That is, the number of false positives is lower compared to false negatives. A possible reason for the situation here is that several weak SEP events have proton fluxes closely fluctuating in the vicinity of the SWPC 'S1' threshold. Detecting such patterns at high precision becomes relevant to reducing FNs.

In Figure 7, we compare the skill scores of STSF with respect to the different prediction windows considered in this work. The average TSS at all prediction windows is $\approx 0.86(\pm 0.02)$. Similarly, the HSS, MCC and GSS scores obtained are also high. There is a slight increment in the scores from T_5 to T_{30}, but reduces at T_{45} and T_{60}. This slight variation is due to subtle changes in the optimal classification threshold $\approx 0.3(\pm 0.03)$ for STSF. The scores for T_{60} only have a marginal reduction ($<2\%$) compared to T_5. Furthermore, there is a decrease of $\approx 7\%(\pm 2\%)$ in skill scores compared to Rotti et al. (2024). This is expected due to extended lead times considered with a larger class imbalance in the data set. Overall, the classification and computational efficiencies we obtain for STSF are highly satisfactory and viable. This means that our work adequately provides a proof of concept of our modeling architecture for short-term predictive capabilities up to a 60-minute window and can be further transformed into implementation for NRT operations.

5. CONCLUSIONS

In this work, we approach the SEP event short-term prediction problem from a time series classification perspective. We primarily use the GSEP data set but extend our negative sample size by including "SEP-quiet" periods. That is, we consider $\geq$C6.0 flares that do not lead to SEP events observed near Earth. Our data set constitutes in-situ time series measurements from the NOAA-GOES missions for solar cycles 22 to 24. We utilize the long band (1–8Å) X-ray measurements from the XRS instrument and proton fluxes (P3, P5, P7) from the SEM instrument onboard GOES missions.

Figure 7. Comparison of skill scores of STSF model on the test set at prediction windows T_5, T_{15}, T_{30}, T_{45}, and T_{60}.

Our data set has a total of 2893 samples, of which 244 are strong SEP events. We defined the target labels based on the SWPC 'S1' threshold for an SEP event. Positive labels are strong events crossing the 10 pfu threshold in the 10 MeV proton channel, and negative otherwise.

Our goal is to develop robust models that can successfully find discriminants between SEP and SEP-quiet patterns. In our recent study (Rotti et al. 2024), we found our methodology of using an ensemble of feature-based univariate time series classifiers to perform very well in classifying between strong and weak SEP events. Specifically, we found the supervised time series forest (STSF; Cabello et al. 2020) classifier to obtain optimal results. In the present work, we implemented our earlier methodology on an extended data set and compared the performances of three models, namely summary statistic classifier (SSC), one-nearest neighbor (1NN) and STSF. Furthermore, we consider a fixed observational window of 11 hours and use multiple prediction windows of prediction interval: 05-, 15-, 30-, 45-, and 60-min to leverage on high confidence for relatively low lead times.

In our model analysis phase, we used the Youden Index (J) to understand the trade-off between the true positive rate (TPR) and the false positive rate (FPR) at different classification thresholds. That is, we estimate an optimal threshold that provides a TPR with an acceptable FPR to make decisions using our classifiers. We considered the validation set to compare the performances of the models based on multiple skill scores, such as TSS, HSS, GSS and MCC. This resulted in STSF performing better compared to SSC and 1NN at all prediction windows. Furthermore, we use an unseen test set to obtain the classification report of the best model, including the contingency tables. For a 60-minute prediction window, we obtain the following scores for STSF: TSS = 0.850; HSS = 0.878; GSS = 0.783; MCC = 0.879. In summary, our results are promising and provide a gateway for our model architecture to be used in an operational context to offer short-term predictions on the arrival of strong SEP events up to 60 minutes. The data set and coding methodology of our model implementation have been made publicly available on Harvard Dataverse: 10.7910/DVN/MWY6H7 under a Creative Commons license.

The applications we foresee for our predictive system to be useful to space agencies across the globe are two-fold. First, the system can serve as a convenient tool for forecasters, for example, at NASA's Space Radiation Analysis Group (SRAG) and NOAA-SWPC in assisting them in broadcasting highly accurate last-minute warnings to astronauts on the surface of the Moon or during extravehicular activities when outside the Earth's magnetosphere. Second, for future interplanetary missions, such as the missions to Mars that NASA is planning, the spacecraft will spend most of its time outside of the Sun-Earth line. Therefore, SEP predictions that apply to the Earth-Moon system are of limited value when the spacecraft-Sun makes a large angle with the Earth-Sun direction. A crewed spacecraft on a journey to Mars would only have to carry a scaled-down version of the GOES instruments to produce short-term warnings of intense solar particle storms. Hence, we expect that our system can be a valuable addition to the tools that space agencies will need to protect astronauts on upcoming lunar and Mars missions.

Software: pandas (McKinney et al. 2010), numpy (Van Der Walt et al. 2011; Harris et al. 2020), sklearn (Pedregosa et al. 2011), sktime (Löning et al. 2019; Löning et al. 2022), matplotlib (Hunter 2007).

APPENDIX

A. CONTINGENCY TABLES

The contingency table elements for SSC and 1NN on the validation set are provided in Tables 4 and 5, respectively, for each of our experiments on changing prediction windows. The values from these tables are used to estimate the skill scores that are defined in Section 4.1 and shown in Figure 6.

Table 4. Contingency tables for the SSC model on the validation set.

| | | $\mathbf{T}_5$ Predicted | | $\mathbf{T}_{15}$ Predicted | | $\mathbf{T}_{30}$ Predicted | | $\mathbf{T}_{45}$ Predicted | | $\mathbf{T}_{60}$ Predicted | |
		Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak
True	Strong	56	24	43	18	42	38	44	36	58	22
	Weak	44	850	37	876	15	879	20	874	45	849

NOTE—Truth tables for the SSC model for different prediction windows. The first column is a shared entry of true labels against predictive labels for each experiment. The elements indicate the number of predictions with respect to the actual occurrences in the validation set.

Prediction windows are shown by a subscript to T. For example, T_5 indicates a prediction window of 5 min.

Model name: SSC - Summary Statistic Classifier

Table 5. Contingency tables for the 1NN model on the validation set.

| | | T$_5$ Predicted | | T$_{15}$ Predicted | | T$_{30}$ Predicted | | T$_{45}$ Predicted | | T$_{60}$ Predicted | |
		Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak	Strong	Weak
True	Strong	44	36	44	36	43	37	43	37	43	37
	Weak	25	869	24	870	25	869	25	869	29	865

NOTE—Truth tables for the 1NN model for different prediction windows. The first column is a shared entry of true labels against predictive labels for each experiment. The elements indicate the number of predictions with respect to the actual occurrences in the validation set.

Prediction windows are shown by a subscript to T. For example, T$_5$ indicates a prediction window of 5 min.

Model name: 1NN - One-nearest Neighbor

REFERENCES

Alberti, T., Laurenza, M., Cliver, E., et al. 2017, The Astrophysical Journal, 838, 59

Aminalragia-Giamini, S., Raptis, S., Anastasiadis, A., et al. 2021, Journal of Space Weather and Space Climate, 11, 59

Anastasiadis, A., Papaioannou, A., Sandberg, I., et al. 2017, Solar Physics, 292, 1

Angryk, R. A., Martens, P. C., Aydin, B., et al. 2020, Scientific data, 7, 1

Arbib, M. A. 2003, The handbook of brain theory and neural networks (MIT press)

Bain, H., Steenburgh, R., Onsager, T., & Stitely, E. 2021, Space Weather, 19, e2020SW002670

Balch, C. C. 2008, Space Weather, 6, S01001, doi: 10.1029/2007SW000337

Beck, P., Latocha, M., Rollet, S., & Stehno, G. 2005, Advances in Space Research, 36, 1627

Bornmann, P. L., Speich, D., Hirman, J., et al. 1996, in GOES-8 and Beyond, ed. E. R. Washwell, Vol. 2812, International Society for Optics and Photonics (SPIE), 291 – 298, doi: 10.1117/12.254076

Boubrahimi, S. F., Aydin, B., Martens, P., & Angryk, R. 2017, in 2017 IEEE International Conference on Big Data (Big Data), IEEE, 2533–2542

Cabello, N., Naghizade, E., Qi, J., & Kulik, L. 2020, in 2020 IEEE International Conference on Data Mining (ICDM), IEEE, 948–953

Camporeale, E. 2019, Space Weather, 17, 1166, doi: 10.1029/2018SW002061

Cane, H., Reames, D., & Von Rosenvinge, T. 1988, Journal of Geophysical Research: Space Physics, 93, 9555

Cane, H. V., McGuire, R. E., & von Rosenvinge, T. T. 1986, ApJ, 301, 448, doi: 10.1086/163913

Chen, P. 2011, Living Reviews in Solar Physics, 8, 1

Cliver, E. W., & D'Huys, E. 2018, The Astrophysical Journal, 864, 48

Cliver, E. W., Schrijver, C. J., Shibata, K., & Usoskin, I. G. 2022, Living Reviews in Solar Physics, 19, 2

Daan, H. 1985, Statistics, and Decision Making in the Atmospheric Sciences, AH Murphy and RW Katz, Eds., Westview Press, 379

Desai, M., & Giacalone, J. 2016, Living Reviews in Solar Physics, 13, 1

Dierckxsens, M., Tziotziou, K., Dalla, S., et al. 2015, Solar Physics, 290, 841

Engell, A., Falconer, D., Schuh, M., Loomis, J., & Bissett, D. 2017, Space Weather, 15, 1321

Falconer, D., Barghouty, A. F., Khazanov, I., & Moore, R. 2011, Space Weather, 9

Faouzi, J. 2022, Machine Learning (Emerging Trends and Applications)

Feynman, J., & Hundhausen, A. J. 1994, Journal of Geophysical Research: Space Physics, 99, 8451

Garcia, H. A. 1994, Solar Physics, 154, 275

Gopalswamy, N., Lara, A., Yashiro, S., Kaiser, M. L., & Howard, R. A. 2001, Journal of Geophysical Research: Space Physics, 106, 29207

Gopalswamy, N., Mäkelä, P., Yashiro, S., et al. 2017, in Journal of Physics Conference Series, Vol. 900, Journal of Physics Conference Series, 012009, doi: 10.1088/1742-6596/900/1/012009

Gopalswamy, N., Yashiro, S., Xie, H., et al. 2008, The Astrophysical Journal, 674, 560

Grubb, R. N. 1975, The SMS/GOES space environment monitor subsystem, NASA STI/Recon Technical Report N

Hansen, L. K., & Salamon, P. 1990, IEEE transactions on pattern analysis and machine intelligence, 12, 993

Harris, C. R., Millman, K. J., Van Der Walt, S. J., et al. 2020, Nature, 585, 357

Heidke, P. 1926, Geografiska Annaler, 8, 301, doi: 10.1080/20014422.1926.11881138

Hosseinzadeh, P., Boubrahimi, S. F., & Hamdi, S. M. 2024, The Astrophysical Journal Supplement Series, 270, 31

Hunter, J. D. 2007, Computing in science & engineering, 9, 90

Jackman, C. H., & McPeters, R. D. 1987, Physica Scripta, 1987, 309

Jiggens, P., Clavie, C., Evans, H., et al. 2019, Space Weather, 17, 99

Kahler, S. W. 1992, ARA&A, 30, 113, doi: 10.1146/annurev.aa.30.090192.000553

Kahler, S. W., Cliver, E. W., & Ling, A. G. 2007, Journal of Atmospheric and Solar-Terrestrial Physics, 69, 43, doi: 10.1016/j.jastp.2006.06.009

Klein, K.-L., & Dalla, S. 2017, Space Science Reviews, 212, 1107

Laurenza, M., Cliver, E., Hewitt, J., et al. 2009, Space Weather, 7

Lavasa, E., Giannopoulos, G., Papaioannou, A., et al. 2021, Solar Physics, 296, 107

Löning, M., Bagnall, A., Ganesh, S., et al. 2019, arXiv preprint, doi: 10.48550/arXiv.1909.07872

Low, B. 1996, Solar Physics, 167, 217

Löning, M., Király, F., Bagnall, T., et al. 2022, sktime/sktime: v0.13.4, Zenodo, doi: 10.5281/zenodo.7117735

Marqué, C., Posner, A., & Klein, K.-L. 2006, The Astrophysical Journal, 642, 1222

McKinney, W., et al. 2010, Data structures for statistical computing in python

Milligan, R. O., & Ireland, J. 2018, SoPh, 293, 18, doi: 10.1007/s11207-017-1233-x

Núñez, M. 2011, Space Weather, 9, S07003, doi: 10.1029/2010SW000640

Onsager, T., Grubb, R., Kunches, J., et al. 1996, in GOES-8 and Beyond, Vol. 2812, International Society for Optics and Photonics, 281–290

Papaioannou, A., Anastasiadis, A., Kouloumvakos, A., et al. 2018, Solar Physics, 293, 100

Parker, E. 1965, Space Science Reviews, 4, 666

Pedregosa, F., Varoquaux, G., Gramfort, A., et al. 2011, Journal of Machine Learning Research, 12, 2825

Posner, A. 2007, Space Weather, 5

Reames, D. V. 1999, Space Science Reviews, 90, 413

—. 2013, Space Science Reviews, 175, 53

Rotti, S., Aydin, B., Georgoulis, M., & Martens, P. 2022, GSEP Dataset, V5, Harvard Dataverse, doi: 10.7910/DVN/DZYLHK

Rotti, S., Aydin, B., Georgoulis, M. K., & Martens, P. C. 2022, ApJS, 262, 29, doi: 10.3847/1538-4365/ac87ac

Rotti, S., & Martens, P. C. 2023, The Astrophysical Journal Supplement Series, 267, 40

Rotti, S. A., Aydin, B., & Martens, P. C. 2024, arXiv e-prints, arXiv:2403.17418, doi: 10.48550/arXiv.2403.17418

Rotti, S. A., Martens, P. C. H., & Aydin, B. 2020, The Astrophysical Journal Supplement Series, 249, 20, doi: 10.3847/1538-4365/ab9a42

Sakoe, H., & Chiba, S. 1978, IEEE Transactions on Acoustics, Speech, and Signal Processing, 26, 43, doi: 10.1109/TASSP.1978.1163055

Sauer, H. H. 1989, SEL monitoring of the earth's energetic particle radiation environment

Schaefer, J. T. 1990, Weather and forecasting, 5, 570

Schapire, R. E. 1990, Machine learning, 5, 197

Schrijver, C. J., & Siscoe, G. L. 2010, Heliophysics: space storms and radiation: causes and effects (Cambridge University Press)

Schwadron, N. A., Townsend, L., Kozarev, K., et al. 2010, Space Weather, 8

Smart, D., & Shea, M. 1992, Advances in Space Research, 12, 303

Swalwell, B., Dalla, S., & Walsh, R. W. 2017, Solar Physics, 292, 1

Usoskin, I. G. 2017, Living Reviews in Solar Physics, 14, 3

Van Der Walt, S., Colbert, S. C., & Varoquaux, G. 2011, Computing in science & engineering, 13, 22

Van Hollebeke, M. A. I., Ma Sung, L. S., & McDonald, F. B. 1975, SoPh, 41, 189, doi: 10.1007/BF00152967

Whitman, K., Egeland, R., Richardson, I. G., et al. 2022, Advances in Space Research, doi: https: //doi.org/10.1016/j.asr.2022.08.006

Woodcock, F. 1976, Monthly Weather Review, 104, 1209 , doi: 10.1175/1520-0493(1976)104⟨1209: TEOYFF⟩2.0.CO;2

Youden, W. J. 1950, Cancer, 3, 32