



# Incorporating Physical Knowledge Into Machine Learning for Planetary Space Physics

Abigail R. Azari<sup>1\*</sup>, Jeffrey W. Lockhart<sup>2</sup>, Michael W. Liemohn<sup>1</sup> and Xianzhe Jia<sup>1</sup>

## OPEN ACCESS

### Edited by:

Bala Poduval,  
University of New Hampshire,  
United States

### Reviewed by:

Simon Wing,  
Johns Hopkins University,  
United States  
Michael S. Kirk,  
The Catholic University of America,  
United States  
Ryan McGranaghan,  
ASTRA, LLC, Spain

### \*Correspondence:

Abigail R. Azari  
azari@umich.edu

### † Present address:

Abigail R. Azari,  
Space Sciences Laboratory, University  
of California, Berkeley, Berkeley, CA,  
United States

### Specialty section:

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

**Received:** 05 March 2020

**Accepted:** 28 May 2020

**Published:** 08 July 2020

### Citation:

Azari AR, Lockhart JW, Liemohn MW  
and Jia X (2020) Incorporating  
Physical Knowledge Into Machine  
Learning for Planetary Space Physics.  
*Front. Astron. Space Sci.* 7:36.  
doi: 10.3389/fspas.2020.00036

<sup>1</sup> Climate and Space Sciences and Engineering Department, University of Michigan, Ann Arbor, MI, United States, <sup>2</sup> Sociology Department, University of Michigan, Ann Arbor, MI, United States

Recent improvements in data collection volume from planetary and space physics missions have allowed the application of novel data science techniques. The Cassini mission for example collected over 600 gigabytes of scientific data from 2004 to 2017. This represents a surge of data on the Saturn system. In comparison, the previous mission to Saturn, Voyager over 20 years earlier, had onboard a ~70 kB 8-track storage ability. Machine learning can help scientists work with data on this larger scale. Unlike many applications of machine learning, a primary use in planetary space physics applications is to infer behavior about the system itself. This raises three concerns: first, the performance of the machine learning model, second, the need for interpretable applications to answer scientific questions, and third, how characteristics of spacecraft data change these applications. In comparison to these concerns, uses of “black box” or un-interpretable machine learning methods tend toward evaluations of performance only either ignoring the underlying physical process or, less often, providing misleading explanations for it. The present work uses Cassini data as a case study as these data are similar to space physics and planetary missions at Earth and other solar system objects. We build off a previous effort applying a semi-supervised physics-based classification of plasma instabilities in Saturn’s magnetic environment, or magnetosphere. We then use this previous effort in comparison to other machine learning classifiers with varying data size access, and physical information access. We show that incorporating knowledge of these orbiting spacecraft data characteristics improves the performance and interpretability of machine learning methods, which is essential for deriving scientific meaning. Building on these findings, we present a framework on incorporating physics knowledge into machine learning problems targeting semi-supervised classification for space physics data in planetary environments. These findings present a path forward for incorporating physical knowledge into space physics and planetary mission data analyses for scientific discovery.

**Keywords:** planetary science, automated event detection, space physics, Saturn, physics-informed machine learning, feature engineering, domain knowledge, interpretable machine learning

## 1. INTRODUCTION

Planetary space physics is a young field for large-scale data collection. At Saturn for example, it was only in 2004 that the first Earth launched object orbited this planet (Cassini) and landed on Titan (Huygens). After arriving Cassini collected data about Saturn and its near-space environment for 13 years, resulting in 635 GB of scientific data (NASA Jet Propulsion Laboratory, 2017a). To put this into perspective, the Voyager I mission which flew by Saturn in 1980 had onboard  $\sim 70$  kB of memory total (NASA Headquarters, 1980). The Cassini mission represents the first large-scale data collection of Saturn. This enabled the field of planetary science to apply statistics to large-scale data sizes, including machine learning, to the most detailed spatio-temporally resolved dataset of the planet and its environment.

This surge of data is not unique to Saturn science. In planetary science broadly, Mars in 2020 has eight active missions roving along the surface and orbiting (Planetary Society, 2020). The Mars Reconnaissance Orbiter alone has already collected over 300 TB of data (NASA Jet Propulsion Laboratory, 2017b). It is commonly accepted that upcoming missions will face similar drastic advances in the collection of scientific data. Traditionally planetary science has employed core scientific methods such as remote observation and theoretical modeling. With the new availability of sampled environments provided by these missions, methods in machine learning offer significant potential advantages. Applying machine learning in planetary space physics differs from other common applications. Cassini's data are characteristic of other planetary and space physics missions like the Magnetospheric Multiscale Mission (MMS) at Earth and the Juno mission to Jupiter. The plasma and magnetic field data collected by these missions are from orbiting spacecraft. This conflates spatial and temporal phenomena. This is a shared characteristic with the broader field of geoscience which often represents complex systems undergoing significant spatio-temporal changes with limitations on quality and resolution (Karpatne et al., 2019).

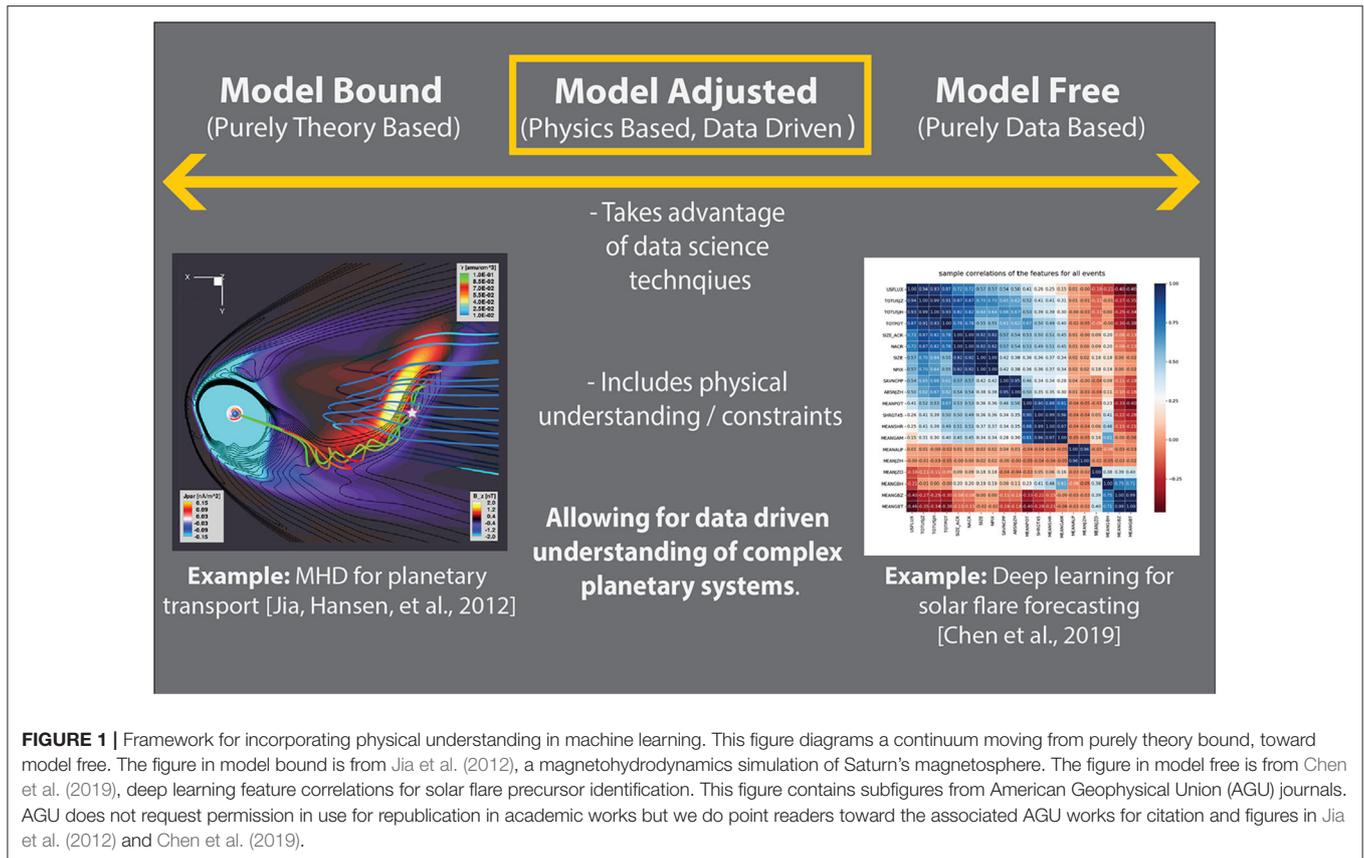
The main use of these data in planetary science is to advance fundamental scientific theories. This requires the ability to infer meaning from applications of statistical methods. Unlike similar missions at Earth, machine learning for space physics data at Saturn has limited direct application to the prediction of space weather. A central interest in space weather prediction is to give lead-time information for operational purposes. As a result, the prediction accuracy in machine learning applications in space weather prediction is seen as paramount. In comparison, at Saturn, machine learning applications require highly interpretable and explainable techniques to investigate scientific questions (Ebert-Uphoff et al., 2019). How to improve machine learning generally from an interpretability standpoint is itself an active research area in domain applications of machine learning (e.g., Molnar, 2019). Within this work we specifically focus on evaluating and implementing interpretable machine learning. Interpretable machine learning usually relies on domain knowledge and is therefore domain specific, but it can be extended to generally refer to models with functional forms simple enough for humans to understand how they make

predictions, such as logical rules or additive factors (Rudin, 2019). Complexity depends in part on what constitutes common knowledge within a domain. Scientists are trained to interpret different models depending on their field. As a result models will range in perceived interpretability across fields. While the final models must be relatively simple in order for humans to understand their decision process, the algorithms which produce optimal interpretable models often require solving computationally hard problems. Importantly, despite widespread myths about performance, interpretable models can often be designed to perform as well as un-interpretable or "black box" models (Rudin, 2019).

In planetary science it's important to discern the workings of a model in order to understand the implications for the workings of physical systems. Interpretability is not the same as explainability: explainability refers to any attempt to explain how a model makes decisions, typically this is done afterwards and without reference to the model's internal workings. Interpretability, however, refers to whether the inner workings of the model, its actual decision process, can be observed and understood (Rudin, 2019). Within this work we are concerned with interpretability in order to gain scientifically actionable results from applied machine learning. The dual challenges of spatio-temporal data and interpretability are compounded for planetary orbiting spacecraft. Complications for orbiting spacecraft can range from rare opportunities for observation, and engineering constraints on spacecraft data transmission. A main interest in this work is to begin to ask: how can machine learning be used within these constraints to answer fundamental scientific questions?

Scientists have approached interpretable machine learning for physics in two ways. First, they have added known physical constraints and relationships into modeling. Within the space weather prediction community, such integration has shown promise in improving the performance of deep learning models over models that do not account for the physics of systems (Swiger et al., 2020). Several fields including biology have argued for an equal value of domain knowledge and machine learning techniques for that reason (see discussion within Coveney et al., 2016). These discussions have culminated in several reviews for scientific fields on the integration of machine learning for data rich discovery (Butler et al., 2018; Bergen et al., 2019). Second, scientists have long tried to use machine learning for the discovery of physical laws from machine learning (e.g., Kokar, 1986). Recently, this work has turned to deep learning tools (e.g., Ren et al., 2018; Raissi et al., 2019; Iten et al., 2020). However, as Rudin (2019) points out, explanations for the patterns deep learning tools find are often inaccurate and at worst, totally unrelated to both the model and the world it models. These two approaches lie on a continuum between valuing increasing data and model freedom, or incorporating physical insight and model constraint.

In **Figure 1**, we present a diagram for considering physical theory and machine learning within the context of theoretical constraints. The examples at one end of the continuum represent applications of traditional space physics from global theory driven modeling, while those at the other end of the continuum



focus on data driven approaches to space weather and solar flare prediction. The model adjusted center presented below takes advantage of data, but limits or constrains the application by merging with domain understanding. Our work is in the middle of the continuum. We leverage domain knowledge about space physics, while also aiming to learn more about the physical system we study. Importantly, we use an interpretable machine learning approach so that we can be more confident in drawing physical insights from the model.

We present comparisons between a range of data sizes and physics incorporation to classify unique plasma transport events around Saturn using the Cassini dataset. As a characteristic data set of space physics and planetary environments, this provides valuable insights toward future implementation of automated detection methods for space physics and machine learning. We focus on three primary guiding axes in this work to address implementations of machine learning. First, we address the performance and accuracy of the application. Second, we consider how to increase interpretability of machine learning applications for planetary space physics. Third, we tackle how characteristics of spacecraft data change considerations of machine learning applications. All of these issues are essential to consider in applications of machine learning to planetary and space physics data for scientific interpretation.

To investigate these questions and provide a path toward application of machine learning to planetary space physics

datasets, we compare and contrast physics-based and non-physics based machine learning applications. In section 2, we discuss the previous development of a physics-based semi-supervised classification from Azari et al. (2018) for the Saturn system within the context of common characteristics of orbiting spacecraft data. We then provide an outline for general physics-informed machine learning for automated detection with space physics datasets in section 3. Section 4 describes the machine learning model set up and datasets that we use to compare and contrast physics-based and non-physics based event detection. Section 5 details the implementation of logistic regression and random forest classification models as compared to this physics-based algorithm with the context of physics-informed or model adjusted machine learning. Section 6 then concludes with paths forward in applications of machine learning for scientific insight in planetary space physics.

## 2. BACKGROUND: SATURN'S SPACE ENVIRONMENT AND DATA

Saturn's near space environment where the magnetic field exerts influence on particles, or magnetosphere, ranges from the planet's upper atmosphere to far from the planet itself. On the dayside the magnetosphere stretches to an average distance of 25 Saturn radii ( $R_S$ ) with a dynamic range between 17 and 29  $R_S$  (Arridge

et al., 2011) ( $1 R_S = 60,268$  km). This distance is dependent on a balance between the internal dynamics of the Saturn system and the Sun's influence from the solar wind. Within this environment a complex system of interaction between a dense disk of neutrals and plasma sourced from a moon of Saturn, Enceladus, interacts with high-energy, less dense plasma from the outer reaches of the magnetosphere (see **Figure 2**).

This interaction, called interchange, is most similar to Rayleigh-Taylor instabilities and results in the injection of high-energy plasma toward the planet. In **Figure 2**, a system of interchange is detailed with a characteristic Cassini orbit cutting through the interchanging region. The red box in this figure is presented as an illustrative slice through the type of data obtained to characterize interchange. One of the major questions in magnetospheric studies is how mass, plasma, and magnetic flux moves around planets. At the gas giant planets of Saturn and Jupiter, interchange is thought to be playing a fundamental role in system-wide transport by bringing in energetic material to subsequently form the energetic populations of the inner magnetosphere, and to transport plasma outwards from the moons. Until Cassini, Saturn never had a spacecraft able to develop statistics based on large-scale data sizes to study this mass transport system.

The major scientific question surrounding studying these interchange injections is what role these injections are playing in the magnetosphere for transport, energization, and loss of

plasma. To answer this question, it's essential to understand where these events are occurring and the dependency of these events on other factors in the system, such as influence from other plasma transport processes and spatio-temporal location. From Cassini's data, several surveys of interchange had been pursued by manual classification, but these surveys disagreed on both the identification of events and resulting conclusions (Chen and Hill, 2008; Chen et al., 2010; DeJong et al., 2010; Müller et al., 2010; Kennelly et al., 2013; Lai et al., 2016). The main science relevant goal was to create a standardized, and automated, method to identify interchange injections. This list needed to be physically justified to allow for subsequent conclusions and comparisons.

In section 2.1, we provide background on the Cassini dataset and summarize the previous development of a physics-based detection method in section 2.2. We then provide a generalized framework in the following section 3 for incorporating physical understanding into machine learning with the development of this previous physics-based method as an example. Subsequent sections investigate comparisons of this previous physics-based effort to other automated identification methods.

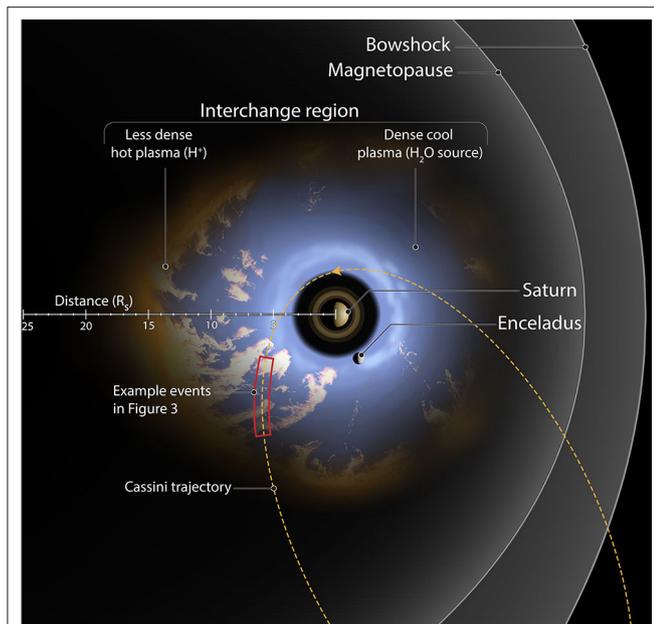
## 2.1. Cassini High-Energy Ion Dataset

Cassini has onboard multiple plasma and wave sensors which are in various ways sensitive to interchange injections. However, none of the previous surveys focused on high-energy ions, which are the primary particle species transported inwards during injections. In **Figure 3**, a series of injections are shown in high-energy (3–220 keV) ions ( $H^+$ ) and magnetic field datasets. This figure shows three large injections between 0400 and 0600 UTC followed by a smaller injection after 0700 most noticeable in the magnetic field data. It is evident from these examples that using different sensors onboard Cassini will result in different identification methods for interchange injections. This was a primary driver in a standardized identification method for these events. The top two panels detail the Cassini Magnetospheric Imaging Instrument: Charge Energy Mass Spectrometer (CHEMS) dataset while the last contains the Cassini magnetometer magnetic field data (Dougherty et al., 2004; Krimigis et al., 2004).

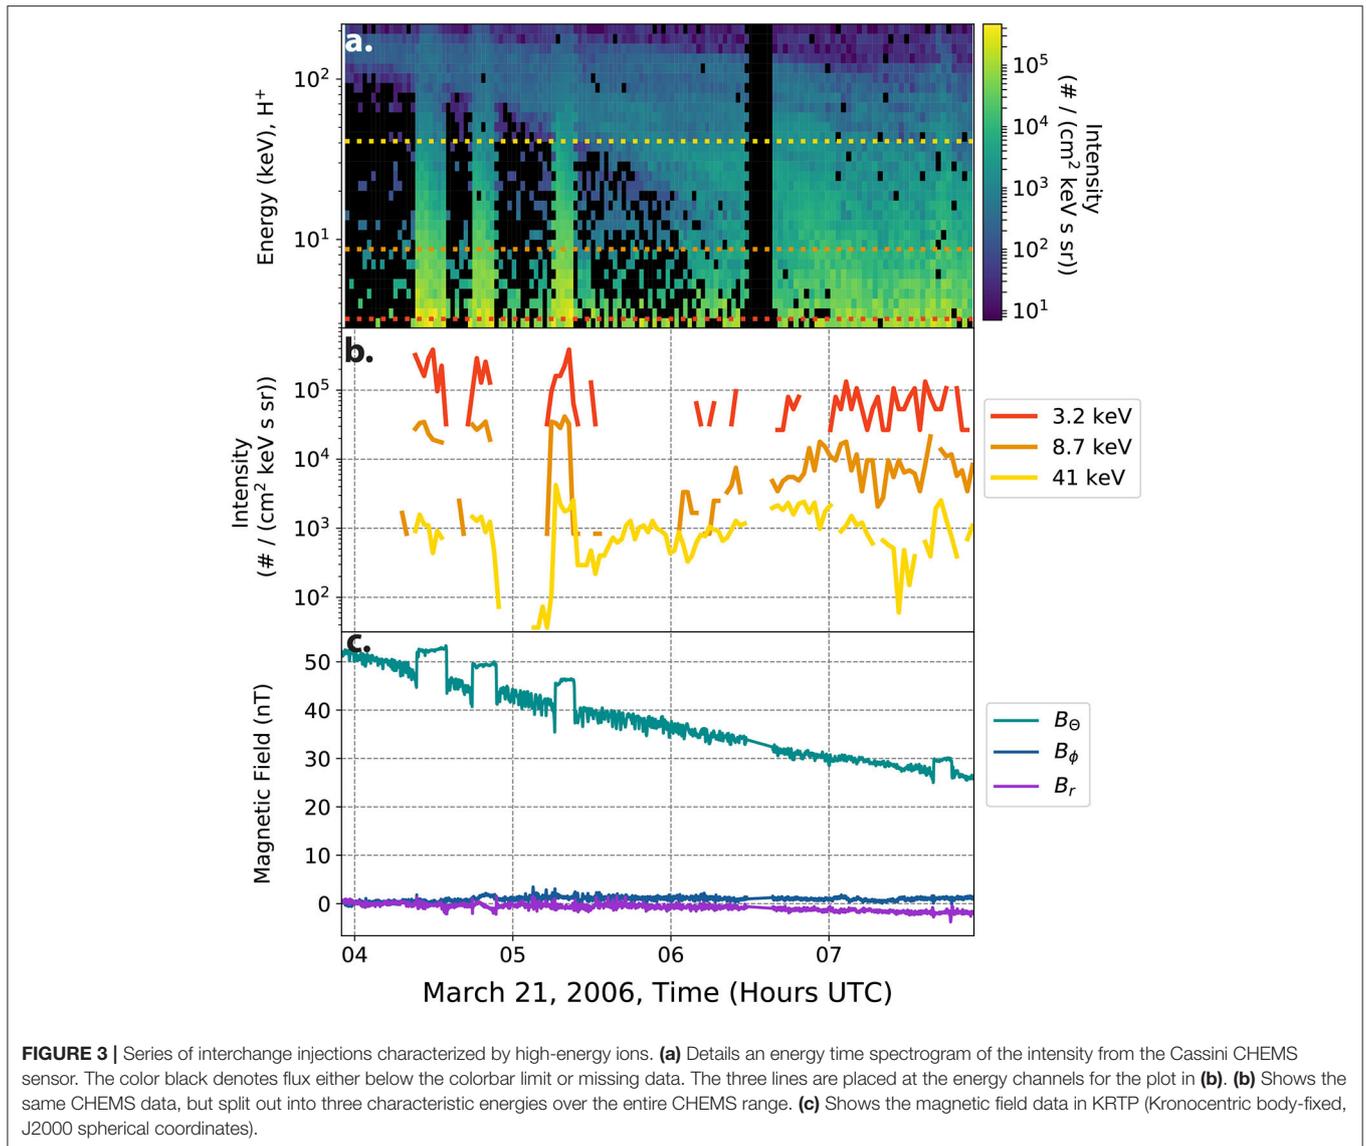
The CHEMS instrument onboard Cassini collected multiple species of ion data and finds the intensity of incoming particles in the keV range of data. This datastream can be thought of as unique energy channels, each with a spacecraft position and time dependence. In **Figure 3b**, three unique energy channels are shown from the overall data in the top panel, to illustrate the nature of these high-energy data. This type of spatio-temporal data is often a characteristic of space physics missions (see Baker et al., 2016, for a review of MMS' data products).

## 2.2. Development of Physics-Based Detection Method

When applying automated or machine learning methods, such data discussed above provides unique challenges and characteristics including: rare events (class imbalances), spatio-temporal sampling, heterogeneity in space and time, extreme high-dimensionality, and missing or uncertain data (Karpatne et al., 2019). These challenges are in addition to desired



**FIGURE 2** | Diagram of interchange injection in the Saturn system. The illustrated orbit is an equatorial Cassini orbit from 2005. Injections are denoted by the pale orange material interspersed with the water sourced plasma from Enceladus. Along the example orbit the red box denotes a hypothetical segment of Cassini data discussed in **Figure 3**. The purpose of developing an automated event detection is to identify the pale orange material traveling toward the planet. This figure is produced in consultation with, and copyright permissions from Falconieri Visuals.



interpretability. It's essential that an interpretable model is used to learn substantive information about this application. One common use of machine learning is to input a large number of variables and/or highly granular raw data (e.g., individual sensor readings or image pixel values) into a model, letting the algorithm sort out relationships among them. Such models are inherently “black boxes” because the number and granularity of variables, not to mention complicated recursive relationships among them, makes it difficult or impossible for humans to interpret (Rudin, 2019). One solution to this issue is to reduce dimensionality to fewer, more meaningful-to-humans inputs. But at the same time, the model needs to be informative, and the inputs need to be meaningful. Incorporating domain knowledge and then letting the model determine their effectiveness in the system of study is a potential framework to consider.

For this reason, when developing a detection method to standardize, characterize, and subsequently build off the detected

list, a physics-based method was chosen to address these unique challenges. This previous effort is discussed in Azari et al. (2018) and the resultant dataset is located on the University of Michigan's Deep Blue Data hub (Azari, 2018). We build on this effort in the present work to provide a new evaluation of alternative solutions for data-driven methods.

To develop this physics-based method, the common problems in space physics data described in Karpatne et al. (2019) were considered and addressed to develop a single dimension array ( $S$ ).  $S$  was then used in a style most similar to a single dimensional logistic regression to find the optimum value for detecting interchange events. This classification was standardized in terms of event severity, as well as physically bounded in definition of events. As a result, it was able to be used to build up a physical understanding of the high-energy dynamics around Saturn's magnetosphere including: to estimate scale sizes (Azari et al., 2018) and to demonstrate the influence of tail

injections as compared to the ionosphere (Azari et al., 2019). Following machine learning practices,  $S$  was designed through cross validation. It was created to perform best at detecting events in a training dataset and then evaluated on a separate test dataset. These sets contain manually identified events and were developed from 10% of the dataset (representing 7,375/68,090 time samples). Training and test dataset selection and limiting spatial selection is of critical importance in spatio-temporal varying datasets. Our particular selection considerations are discussed in following sections. The training set was used to optimize the final form of  $S$ . The test dataset was used to compare performance and prevent over fitting. The same test and training datasets are used in the following sections.

$S$  was developed in Azari et al. (2018) to provide a single-dimension parameter which separated out the multiple dependencies of energy range and space while dealing with common challenges in space physics and planetary datasets.  $S$  is calculated from  $S_r$  by removing the radial dependence through normalization. In mathematical form,  $S_r$  can be written as:

$$S_r = \sum_{e=0}^{14} w^{(Z_{e,r}-C)} \quad (1)$$

$S$  can be thought of as a single number which describes the intensification of particle flux over a normalized background. In other words,  $S$  can be calculated as:  $S = (S_r - \bar{S}_r)/\sigma_{S_r}$ . In which  $\bar{S}_r$  is the average radially dependent average and  $\sigma_{S_r}$  the radially dependent standard deviation. These calculations allow for  $S$  to be used across the entire radial and energy range for optimization in units of standard deviation. The variables  $w$  and  $C$  represent weighting values which are optimized for and discussed in the following section. The notations of  $e$  and  $r$  represent energy channel and radial value.  $Z_{e,r}$  represents a normalized intensity value observed by CHEMS. This is similar to the calculation of  $S$  from  $S_r$ .

Additional details on the development, and rationale behind,  $S$  are described in section 3 as a specific example for a general framework for inclusion of physical information into machine learning.

The final form of  $S$  depends non-linearly on the intensity values of the CHEMS sensor and radial distance. In Figure 4, we show the dependence of the finalized  $S$  value over the test dataset for the intensity at a single energy value of 8.7 keV and over all radial distances. Within this figure the events in the test dataset are denoted with dark pink dots. From Figures 4d,e, it's evident that  $S$  disambiguates events from underlying distributions, for example in Figure 4b. By creating  $S$  it was possible to create a single summary statistic which separated events from a background population.

The strategies pursued in developing  $S$  are most applicable for semi-supervised event detection with space physics data. They can, however, prove a useful guide in starting to incorporate physical knowledge into other applications in heliophysics and space physics. Within the previous effort we used the model optimization process from machine learning to guide a physics incorporated human effort. This was a solution to

incorporating the computational methods employed in machine learning optimization to a human-built model. The end result was optimized in a similar fashion as machine learning models but through manual effort to ensure physical information preservation. Moving from this effort, we now present a framework for expanding the style of integrating human effort and physical-information into other applications for space physics data.

Below we provide a framework for incorporating physical-understanding into machine learning. In each strategy we discuss common issues in space physics data, using a similar phraseology as Karpatne et al. (2019). In addition to characteristics in the structure of geoscience data, we also add interpretability as a necessary condition. For space physics and planetary data, the challenges within Karpatne et al. (2019) are often compounded and where appropriate we note potential overlap. After each strategy, we provide a walk-through of the development of  $S$  employed in Azari et al. (2018).

### 3. FRAMEWORK FOR PHYSICS INCORPORATION INTO MACHINE LEARNING

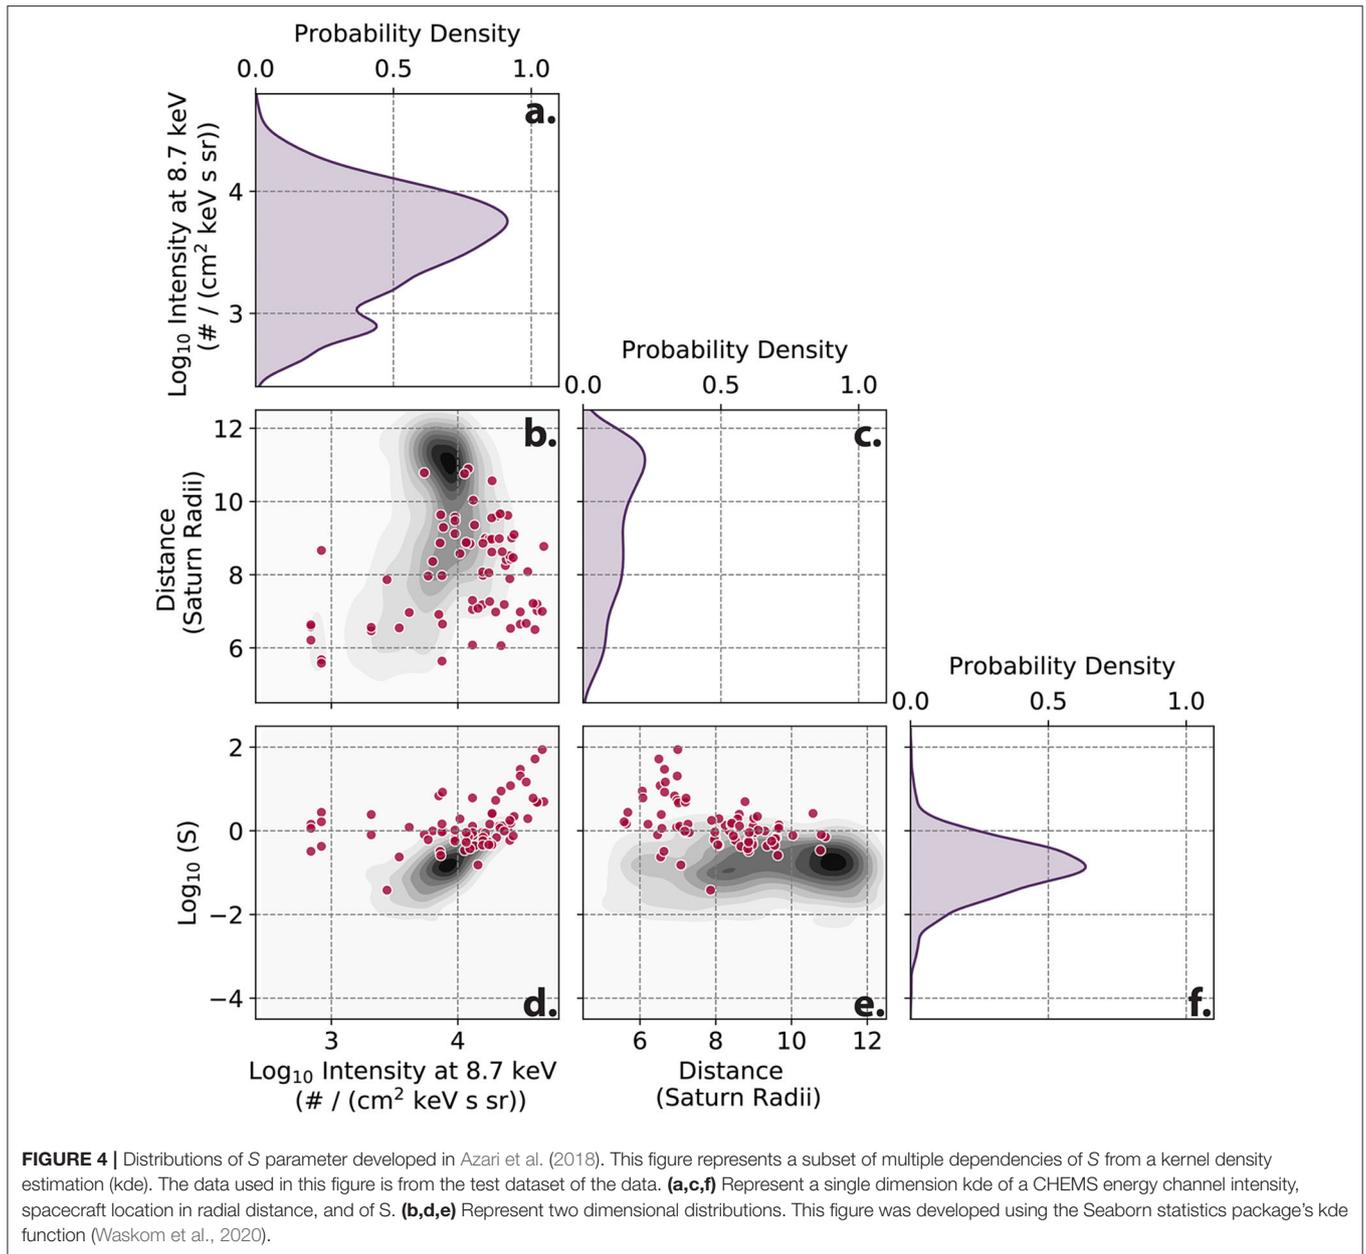
This framework focuses on interpretable semi-supervised event detection with space physics data from orbiters for the end goal of scientific analysis. Depending on the problem posed certain solutions could be undesirable. For a similarly detailed discussion on creating a machine learning workflow applied to problems in space weather (see Camporeale, 2019). The framework presented here can be thought of as a directed application of feature engineering for space physics problems, mostly for requiring interpretability. In general the strategies below provide a context for careful consideration of the nature of domain application which is essential for applications of machine learning models to gather scientific insights.

1. **Limit to region of interest.** Orbiting missions often range over many environments and limiting focus to regions of interest can assist in automated detection by increasing the likelihood of detection of events.

**Issues:** heterogeneity in space and time, rare events (class imbalance)

**Example:** The Cassini dataset represents a wide range of sampled environments, the majority of which do not exhibit interchange. In addition, the system itself undergoes seasonal cycles, changing in time, presenting a challenge to any long-ranging spatial or temporal automated detection. The original work targeted a specific radial region between 5 and 12  $R_S$  in the equatorial plane. This region is known to be sensitive to interchange from previous studies. Similarly, each season of Saturn was treated to a separate calculation of  $S$ , allowing for potential temporal changes to the detection of interchange.

2. **Careful consideration of training and test datasets.** Due to the orbiting nature of spacecraft, ensuring randomness in training and test datasets is usually not sufficient to create a representative set of data across space and time. For event studies, considerations of independence for training and test



dataset while containing prior and post-event data (at times critical for event identification) are important. This is similar to recent strides in activity recognition studies with spatio-temporal data, in which training set considerations drastically affect the accuracy of activity classification (e.g., Lockhart and Weiss, 2014a,b).

**Issues:** heterogeneity in space and time, spatio-temporal data, rare events, small sample sizes

**Example:** While the test and training set represent 10% of the data for the worked example, the 10% was taken such that it covered the widest range of azimuthal and radial values, while still being continuous in time and containing a range of events.

3. **Normalize and/or transform.** Many space environments have a spatio-temporal dependent background. Normalizing separately to spatial or other variables will address these dependencies and can prove advantageous if these are not critical to the problem.

**Issues:** heterogeneity in space and time, spatio-temporal sampling, multi-dimensional data

**Example:** As seen **Figure 4b** flux values depend on radial distance and energy value. Similarly, flux exhibits log scaling, where values can range over multiple powers of 10 in the span of minutes to hours as seen in **Figure 3**. To handle the wide range of values from the CHEMS sensor, each separate energy channel's intensity was first converted into logarithmic

space before then being normalized by subtracting off the mean and dividing by its standard deviation. Effectively, this transforms the range of intensities to a near-normal distribution dependent on radial distance and energy value (see  $Z_{e,r}$  in Equation 1). A similar treatment is performed on creating the final  $S$  from  $S_r$ . This is important due to the commonality of normalcy assumptions in which models can assume normally distributed data on the same scale across inputs.

- 4. Incorporate physical calculations.** Space physics data can come with hundreds if not thousands of features. While many machine learning techniques are designed for just this kind of data, they do not typically yield results that are amenable to human interpretation and scientific insight into the processes of physical systems. They express a complex array of relationships among raw measurements that do little to help humans build theory or understanding. Summary statistics like summing over multiple variables, or taking integrals, can preserve a large amount of information from the raw data for the algorithm while leaving scientists with smaller sets of relationships between more meaningful variables to interpret. For other fields rich in noisy and incomplete time-series data with a longer history of automated detection methods, summary statistic transformations have been a valuable way of handling this type of data for improved performance (e.g., Lockhart and Weiss, 2014a).

**Issues:** interpretability, multi-dimensional data, missing data

**Example:** To address missing values . building up summary statistics, for example through summing over multiple energy channels can help. This creates an particle pressure like calculation (see sum in Equation 1). Particle pressure itself is not used to identify events, as the ability to tune the exact parameters was desired in the identification of injections and developing  $S$  proved more reliable. This allows for the lower 14 energy channels to contribute without removing entire timepoints from the calculation where partial data is missing and also increasing interpretability of the end result . Only the lower 14 channels are used as the higher energy channels also show long duration background from earlier events drifting in the Saturn environment (see **Figure 3**).

- 5. Compare with alternate metrics.** Dependent on your use case, the trade-off costs between false positives and false negatives could be different from the default settings in standard machine learning tools. Investigating alternate metrics of model performance and accuracy are useful toward increasing interpretability.

**Issues:** interpretability, rare events (class imbalance)

**Example:** In the training and test datasets only 2.4% of the data exist in an event state. This proves to be challenging for then finding optimum detection due to the amount of false positives and usage for later analysis. In Equation (1) scaling factors of  $w$  and  $C$  are introduced. These scale factors are chosen by optimizing for the best performance of the Heidke Skill Score (HSS) (Heidke, 1926). HSS is more commonly used in weather forecasting than in machine learning penalty calculations but has shown potential for handling rare events (see Manzato, 2005, for a discussion of HSS). In section

5, we evaluate how HSS performs as compared to other regularization schemes (final values:  $w = 10$ ,  $C = 2$ ).

- 6. Compare definitions of events, consider grounding in physical calculations.** Much of the purpose of developing an automated detection is to standardize event definitions. Developing a list of events then can become tricky.

**Issues:** lack of ground truth, interpretability, rare events (class imbalance)

**Example:** At this point in the calculation of  $S$ , there is a single number, in units of standard deviations, for each time point. This calculation so far, takes in the flux of the lowest 14 energy channels of CHEMS before normalizing and combining these values to return a single value at each time. This number is higher (in the useful units of standard deviation) for higher flux intensifications and lower for flux drop outs. The final question becomes at which  $S$  value should an event be considered real or false.

Based on the training dataset, 0.9 standard deviations above the mean of  $S$  is the optimum parameter for peak HSS performance. As discussed in section 2.2 0.9 was determined through optimizing against the training set. Since  $S$  is in terms of standard deviations, additional higher thresholds can be implemented to sub-classify events into more or less severe cases with a physical meaning (ranking). This allowed for the application as a definition task with a physical justification.

- 7. Investigate a range of machine learning models and datasets.** Incorporating a range of machine learning models, from the most simple to the most complex in addition to varying datasets, can offer insights in the nature of the underlying physical data.

**Issues:** interpretability

**Example:** In developing  $S$ , alternative feature inclusions were considered.  $S$  was settled on for its grounding in physical meaning. A secondary major consideration was its accuracy compared to other machine learning applications. In the following sections we discuss additional models.

As similarly discussed within Camporeale (2019), the desire to incorporate physical calculations comes from an interest in using machine learning for knowledge discovery. In the use cases of interest here, both the needs for accuracy and interpretability are essential. These presented strategies are designed to improve the potential performance for semi-supervised classification problems and the interpretability for subsequent physical understanding. Creating the final form of  $S$  was a labor intensive process to create and then optimize. Due to  $S$ 's non-linear dependence on the features shown in **Figure 4**, this was a non-trivial task. Similarly expanding  $S$  into additional dimensions is challenging. This is where the machine learning infrastructure offers significant advantages as compared to the previous effort. In the following sections 4 and 5 we discuss alternative solutions to identification of interchange.

## 4. METHODS: MODELS AND EXPERIMENTAL SETUP

In the previous physics-based approach, events were defined through intensifications of  $H^+$  only, allowing for comparisons

to other surveys and advancement of the understanding of events. This was a non-intuitive approach as common logic in application of machine learning algorithms suggests that greater data sizes will result in additional accuracy given a well-posed problem. To explore both the potential for higher accuracy as well as interpretability of the application, we compare the performance of two distinct machine learning models with access to varying data set sizes. Below we discuss models we use in this comparison effort.

#### 4.1. Models

Two commonly used machine learning models for supervised classification are logistic regression and random forest classification. Both are considered standard classification models when applying machine learning and performing comparative studies (Couronné et al., 2018). While both models can be interpreted by humans, the additive functional form of logistic regression and the broad literature on interpreting it make it highly interpretable. Random Forest models consist of easy to interpret logical rules, but the large numbers and weighted combinations of those rules mean it is less interpretable (Rudin, 2019). The original physics-based algorithm was designed with a logistic regression method in mind, but with significant adjustment. Comparisons to this model are directly informative as a result. Logistic regression categorizes for binary decisions by fitting a logistic form, or a sigmoid. Logistic regression is a simple, but powerful, method toward predicting categorical outcomes from complex datasets. The basis of logistic regression is associated with progress made in the nineteenth century in studying chemical reactions, before becoming popularized in the 1940s by Berkson (1944) (see Cramer, 2002, for a review). When implemented and optimized using domain knowledge, highly interpretable models, like logistic regression, generally perform as well as less interpretable models and even deep learning approaches (Rudin, 2019).

Random forest in comparison classifies by building up collection of decision trees trained on random subsets of the input variables. The predictions of all trees are then combined in an ensemble to develop the final prediction. Similar to logistic regression, the method of random forest has been built over time with the most modern development associated with Breiman (2001). While logistic regression requires researchers to specify the functional form of relationships among variables, random forests add complexity toward classification decisions, by allowing for arbitrary, unspecified non-linear dependencies between features, also known as model inputs.

The models used within this chapter are from the scikit-learn machine learning package in Python (Pedregosa et al., 2011). Within the logistic regression the L2 (least squares) regularization penalty is applied. Within the random forest a grid search with 5-fold cross-validation is used to find the optimum depth between 2 and 5, while the number of trees is kept at 50. These search parameters are chosen to constrain the random forest within the perspective of the noisy nature of the CHEMS dataset and to prevent over fitting. Alterations to this tuning parameter scheme are not seen to alter the results in the following section. Events are relatively rare in the data (2.4% of the data in the training

and test datasets corresponds with an event), and this can bias the fit of models. As such, unless otherwise noted, we use class weighting to adjust the importance of data from each class (event and non-event) inversely proportional to its frequency so that the classes exert balanced influence during model fitting. This results in events weighted higher more important than non-events due to their rarity. Performance is shown in section 5 against the test dataset defined above.

#### 4.2. Dataset Definitions and Sizes

To explore the performance of logistic regression and random forest, four distinct subsets of the Cassini plasma and magnetic field data are utilized ranging in data complexity and size as follows:

1. S\C (Spacecraft) Location and Magnetic Field  
6 features, 68,090 time samples
2. S\C Location, Magnetic Field, and H<sup>+</sup> flux (3–220 keV)  
38 features, 68,090 time samples
3. Low Energy H<sup>+</sup> flux (3–22 keV)  
14 features, 68,090 time samples
4. Azari et al. (2018) (S Value)  
1 feature, 68,090 time samples

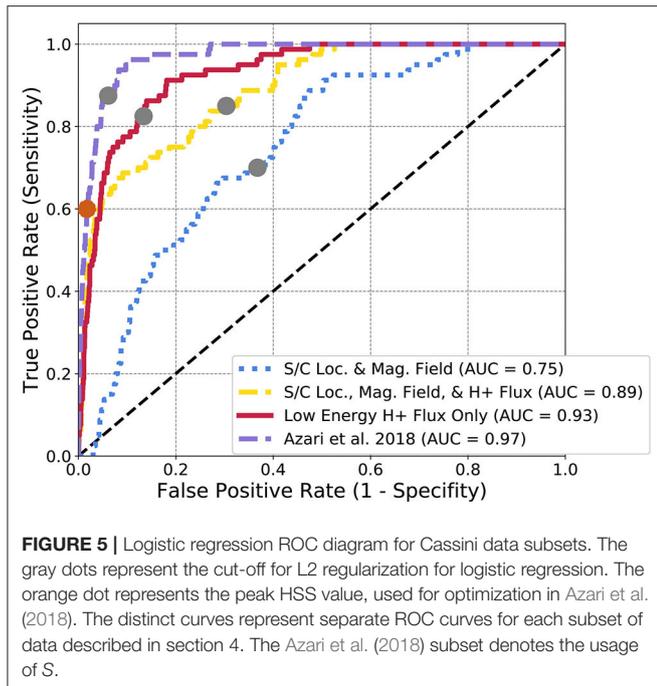
These subsets are chosen to represent additional features, complexity, and physics inclusion. All of these subsets should be sensitive in varying amounts toward identification of interchange injections as evidenced in **Figures 3, 4**. The first two datasets are a comparison of increasing features that should assist in identification of interchange injection. The third dataset includes less features, but is the originator most similar to the derived parameter from Azari et al. (2018). The final dataset contains the single summary statistic array of the S parameter. In the following result section, these four dataset segments are used to evaluate the two models.

### 5. RESULTS AND DISCUSSION

We are interested in evaluating how the former physics-based S parameter performs with other commonly used subsets of space physics data. Our primary goal in this section is to investigate the trade off between the performance of these more traditional models and their interpretability, and therefore usage for scientific analyses. We complete this through applying supervised classification models and evaluate the ease of interpretability and their relative performance.

#### 5.1. Supervised Logistic Regression Classification

In **Figure 5**, the ROC curve of a logistic regression for all four subsets of Cassini data is presented. ROC or receiver operating characteristics, are a common method employed for visualizing the efficacy of classification methods (see Fawcett, 2006, for a generalized review of ROC analysis). ROC curves in this particular example are created by sweeping over a series of classification thresholds. Ideally a perfect classifier will result in a curve that carves a path nearest to the upper left corner.



Area under the curve, or AUC is presented as a metric to understand the overall performance of each logistic regression evaluation. AUC has the ideal parameters of ranging between 0 and 1, with 0.5 representative of random guessing, 1 representing perfect classification, and 0 as the inverse of truth. AUC can be thought of as an average accuracy of a model and isn't sensitive to class-balance and thresholds. ROC curves present the ratios of true positive rate (y-axis) to false positive rate (x-axis). This can be thought of as the trade off for classifiers between events successfully identified (y-axis), and events unsuccessfully identified (x-axis).

The purple curve represents the logistic regression with only the derived physics-based  $S$  as an input. This is rather redundant with optimizing by hand as it's a single variable space. Instead the purple curve is provided as a benchmark against the identical performance and curves found within Azari et al. (2018). From this figure, this single summary statistic ( $S$ ) outperforms all other subsets of Cassini data with an AUC approaching near 1.0 (0.97). This is evidence for the current case, that incorporating physical information, even at the expense of greater dataset size improved the performance of certain machine learning applications.

Following this it is not the largest dataset that has the second best performance. Instead, the red curve which contains only the low energy  $H^+$  intensities shows the best performance of the non physics-adjusted datasets. The magnetic field is a useful parameter for the prediction of interchange as demonstrated in Figure 3, but the form of the logistic regression is unable to use this information successfully. This is possibly due to the higher time resolution needed for interchange identification from magnetic field data and any future identification work needs to focus on adjusting the magnetic field inputs and models. The current dataset is processed such that each time point

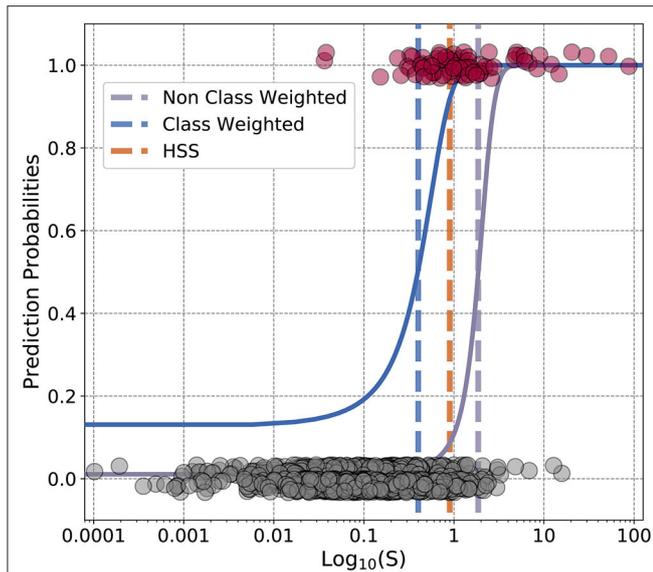
in the CHEMS set is matched with a single magnetic field vector. Normally within interchange analyses, the magnetic field information is of a much higher resolution. It is likely if a study pursued solely magnetic field data of higher time resolution and processed these data to represent pre and post event states dependent on time, the performance of the magnetic field data would be improved. It's evident from Figure 4, that  $S$  exhibits non-linear behavior from the distribution of  $S$  on intensity, distance, and energy. Similarly the magnetic field values likely range over a far range due to the background values, that the linear dependency requirements of logistic regression are unable to use this information. Without the flux data especially (the blue curve) logistic regression is unable to predict interchange as compared to the previous physics-based parameter.

The AUC doesn't capture the entire picture for our interest. While it shows the performance of the algorithm, it contains information for multiple final classifications of events. The gray dots on Figure 5 demonstrates the chosen cut-point for L2 regularization for class weighted events, or the final classification decision for an optimal trade between real events and false events. Within the previous section, the Heidke Skill Score or HSS was discussed as the final threshold separating events from non-events (denoted as the orange dot on Figure 5). Deciding the threshold of what separates an injection event from a non-event is critical for the implementation of statistical analysis on the results especially in this case, in which non-events outnumber events at a ratio of  $\sim 50:1$ . One solution would be to rank events, in similar style of the previous work of  $S$  with categories of events (Azari et al., 2018).

## 5.2. Rare Event Considerations

We now move to evaluating the previous HSS optimization to the logistic regression L2 regulation for both class weighted and non-class weighted models. In Figure 6, the final forms of the weighted and non-weighted logistic regression for the trivial 1 dimensional array case of the  $S$  parameter are shown. The thresholds for the final decisions and for HSS are shown as vertical lines (the orange dashed line represents HSS). Due to the extreme imbalance of non-events to events, implementing class weighting results in large shifts between what is considered an injection event or not. We suggest that the class imbalance inherent in this problem is the main rationale between the differences of HSS and other regularizations. Between the two decision points of the blue and purple vertical lines there are 46 real events, but 202 non-events. This means that if using class-weighting in logistic regression for this problem, 202 non-events would be classified as events. Non-intuitively, for this application where the final events are used to understand the Saturn system, it's advantageous to use a non-class weighted model, as it limits the non-events. However the un-class-weighted model results in removing many real events as well as can be seen in the bulk of the pink events (real events) being misclassified by the purple vertical line.

The Heidke Skill Score provides an in-between choice of these by providing a higher threshold than the class-weighted, and lower than non-class weighted. The logistic regression for the  $S$  parameter shown here is easily intuited since the X-axis

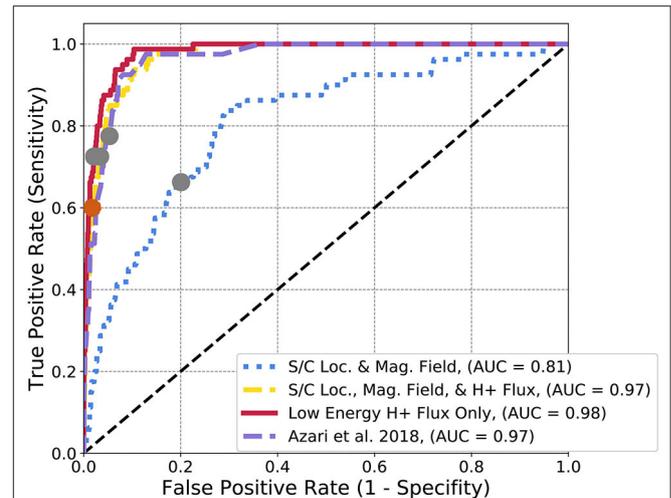


**FIGURE 6** | Finalized logistic regression against test dataset. The gray dots represent the test dataset values of non-events, and the pink of events. The scatter in the dots around 0 and 1 are for aesthetic reasons and do not represent offset values. This figure contains logistic regressions performed on the physics-based parameter from Azari et al. (2018). The blue curve represents a class-weighted model and the purple without class weights. Similarly the dashed lines for blue and purple represent the finalized cut-off points for the class-weighted and un-weighted models. The orange dashed line represents the HSS optimization used within Azari et al. (2018). The x-axis is in logarithmic scale to demonstrate the range of the values,  $S$  itself does span both negative and positive values. From being presented in logarithmic space this gives the false illusion that the blue curve does not approach zero.

represents only one variable. The power of machine learning however is most advantageous in multiple dimensions. HSS has shown to be a more applicable metric for rare events. Other skill scores, such as the True Skill Score have also shown promise in machine learning applications to space physics (Bobra and Couvidat, 2015). Skill score metrics themselves have a long and rich history in space physics before more recent applications in machine learning with interest originating in space weather prediction (see Morley, 2020, for an overarching review of space weather prediction). We also direct the reader to discussions of metrics for physical model and machine learning prediction of space weather (Liemohn et al., 2018; Camporeale, 2019). How can these traditional metrics for space applications be integrated into the regularization schemes? Future work in machine learning applications should consider shared developments between the physical sciences communities usage of skill scores and regularization of models.

### 5.3. Supervised Random Forest Classification

In Figure 7, the ROC diagram for the same subsets of data but for a random forest model are presented. In this case, unlike the logistic regression, other subsets of data can reproduce the same performance (or AUC) as the derived parameter.



**FIGURE 7** | Random forest ROC diagram for Cassini data subsets. The gray dots represent the final optimization location for random forest classification. The orange dot represents the peak HSS value, used for optimization in Azari et al. (2018). The distinct curves represent separate ROC curves for each subset of data described in section 4. The Azari et al. (2018) subset denotes the usage of  $S$ .

All curves, with the exception of the spacecraft location and magnetic field, quickly approach or slightly surpass the AUC of the physics-based parameter at 0.97, with small differences in the performance of the low energy  $H^+$  flux (0.98) and of the combined spacecraft location, all flux, and magnetic field (0.97). The model form of random forest allows for non-linear behavior in the intensity and magnetic field data to find injection events. Increasing the features then helps in the case of random forest whereas it did not for logistic regression. Similar to the logistic regression, HSS results in a different ratio between true positive rate and false positive rate than the random forest model cut-off point with the gray dots.

Comparing back to logistic regression, even with a relatively complex model such as random forest, the AUC of the best ROC curves are near-identical. Given that  $S$  is an array, this is not that surprising. In both cases the physics-derived parameter outperforms or is effectively equivalent to all other data subsets, including those with access to a much richer information set and therefore more complex model. For the application of interpretability for then gathering scientific conclusions, logistic regression is advantageous as it presents a much simpler model. However, random forest, has shown ability to mimic the underlying physics adjustments through selection of datasets.

Within these results, it's evident that the  $S$  parameter performs as well as simplistic machine learning models. Given that  $S$  is also grounded in a physics-based definition dependent on solely a variable flux background, this offers advantages to subsequent usage in scientific results. However, many of the adjustments in creating  $S$  can be implemented into other space physics data, and integrated into machine learning as evidenced here. In the description of the development of  $S$ , several challenges in geoscience data from the framework discussed in Karpatne

et al. (2019), and CHEMS specific solutions were presented. From the above evaluation, it is evident that applications of machine learning are useful to the task of automated event detection from flux data, but with diminishing interpretability. A potential solution to both enhancing the interpretability, similar to the  $S$  based parameter, but also incorporating the advantages of machine learning is presented in **Figure 1**. Rather than consider incorporation of physics-based information as deleterious to the implementation of machine learning, we have found that including this information simplifies the application, enhances the interpretability, and improves the overall performance.

## 6. CONCLUSION AND FUTURE DIRECTIONS

Planetary space physics has reached a data volume capacity at which implementation of statistics including machine learning is a natural extension of scientific investigation. Within this work we addressed how machine learning can be used within the constraints of common characteristics of space physics data to investigate scientific questions. Care should be taken when applying automated methods to planetary science data due to the unique challenges in spatio-temporal nature. Such challenges have been broadly discussed for geoscience data by Karpatne et al. (2019), but until now limited attention in comparison to other fields has been given toward reviews of planetary data.

Within this work we have posed three framing concerns for applications of machine learning to planetary data. First, it's important to consider the performance and accuracy of the application. Second, it's necessary to increase interpretability of machine learning applications for planetary space physics. Third, it's essential to consider how the underlying issue characteristics of spacecraft data changes applications of machine learning. We argue that by including physics-based information into machine learning models, all three concerns of these applications can be addressed.

For certain machine learning models the performance can be enhanced but importantly in this application, the interpretability improves along with handling of characteristic data challenges. To reach this conclusion we presented a framework for incorporating physical information into machine learning. This framework targeted considerations for increasing interpretability and addressing aspects of spacecraft data into machine learning with space physics data. In particular, it addresses challenges such as the spatio-temporal nature of orbiting spacecraft, and other common geoscience data challenges (see Karpatne et al., 2019). After which we then cross-compared a previous physics-based method developed using the strategies in the framework to less physics-informed but feature rich datasets.

The physics-based semi-supervised classification method was built on high-energy flux data from the Cassini spacecraft to Saturn (see Azari et al., 2018). In investigating the

accuracy of machine learning applications, we demonstrated this physics-based approach outperformed automated event detection for simple logistic regression models. It was found that traditional regularization through L2 penalties both under, and overestimated ideal cutoff points for final event classification (depending on class weighting). Instead, metrics more commonly used in weather prediction, such as the Heidke Skill Score, showed promise in class imbalance problems. This is similar to work demonstrating the applicability of True Skill Score in heliophysics applications (Bobra and Couvidat, 2015). Future work should consider building on the rich history of prediction metrics in the space physics community for shared development between the physical sciences usage of skill scores and in regularization of models.

While logistic regression is a more interpretable model, random forest proved that with the addition of more and lower level variables from the Cassini mission, the model could approximate our physics-based logistic model successfully. In this case physics-informed or model adjusted machine learning, can each the same performance but with different levels of interpretability, thus different ability to draw further conclusions about implications of the results. The logistic approach provides a coefficient and threshold for a meaningful physical quantity,  $S$ , effectively the normalized intensification of particle flux. The random forest approach can provide an "importance" score for  $S$  or show a large number of conjunction rules involving it, but neither is as useful for human analysts. A forest model using a large number of raw variables instead of a small number of more meaningful ones like  $S$  is even harder for humans to make sense of. Deep neural networks, as multi-layered webs of weighted many-to-many relationships, are even less informative for human analysts interested in understanding the workings of the model and physical system. Further, findings that the interpretable model performs as well or better than other approaches demonstrate that, despite the widespread myth to the contrary, there is no inherent tradeoff between performance and interpretability (Rudin, 2019). For example, the ability to further split and define identified events based on their flux intensity using  $S$  gives the ability to address further scientific questions as to the fundamental mechanisms behind the interchange instability itself. The simplistic model of logistic regression which results in the same performance as random forest is highly advantageous for the current task.

The framework and comparison presented here opens up avenues toward consideration of applying machine learning to answer planetary and space physics questions. In the future, cross-disciplinary work would greatly advance the state of these applications. Particularly within the context of interpretability toward scientific conclusions through physics-informed, or model adjusted machine learning. The inclusion of planetary science and space physics domain knowledge in application of data science allows for the pursuit of fundamental questions. We have found that incorporating physics-based information increases the interpretability, and improves the overall performance of machine learning applications for scientific insight.

## DATA AVAILABILITY STATEMENT

The events analyzed for this study can be found in the Deep Blue Dataset under doi: 10.7302/Z2WM1BMN (Azari, 2018). The original datasets from the CHEMS (Krimigis et al., 2004) and MAG (Dougherty et al., 2004) instruments can be found on the NASA Planetary Data System (PDS). Details on the most recent datasets for CHEMS and MAG can be found on the Cassini-Huygens Archive page at the PDS Planetary Plasma Interactions node (<https://pds-ppi.igpp.ucla.edu/mission/Cassini-Huygens>). Associated data not included in the above repositories can be obtained through contacting the corresponding author.

## AUTHOR CONTRIBUTIONS

We use the CRediT (Contributor Roles Taxonomy) categories for providing the following contribution description (see Brand et al., 2015). AA led the conceptualization and implemented the research for this manuscript including the investigation, visualization, formal analysis, and original drafting of this work. JL assisted in the conceptualization and discussions of methodology in this work along with editing the manuscript. ML provided funding acquisition, resources, supervision, and assisted in conceptualization along with editing the manuscript. XJ provided funding acquisition, resources, supervision, and assisted in conceptualization. All authors contributed to the article and approved the submitted version.

## REFERENCES

- Arridge, C. S., André, N., McAndrews, H. J., Bunce, E. J., Burger, M. H., Hansen, K. C., et al. (2011). Mapping magnetospheric equatorial regions at Saturn from Cassini Prime Mission observations. *Space Sci. Rev.* 164, 1–83. doi: 10.1007/s11214-011-9850-4
- Azari, A. R. (2018). *Event List for “Interchange Injections at Saturn: Statistical Survey of Energetic H<sup>+</sup> Sudden Flux Intensifications”*. University of Michigan - Deep Blue. doi: 10.7302/Z2WM1BMN
- Azari, A. R. (2020). *A data-driven understanding of plasma transport in Saturn’s magnetic environment*. (Ph.D. thesis). University of Michigan, Ann Arbor, MI, United States.
- Azari, A. R., Jia, X., Liemohn, M. W., Hospodarsky, G. B., Provan, G., Ye, S.-Y., et al. (2019). Are Saturn’s interchange injections organized by rotational longitude? *J. Geophys. Res.* 124, 1806–1822. doi: 10.1029/2018JA026196
- Azari, A. R., Liemohn, M. W., Jia, X., Thomsen, M. F., Mitchell, D. G., Sergis, N., et al. (2018). Interchange injections at Saturn: Statistical survey of energetic H<sup>+</sup> sudden flux intensifications. *J. Geophys. Res.* 123, 4692–4711. doi: 10.1029/2018JA025391
- Baker, D. N., Riesberg, L., Pankratz, C. K., Panneton, R. S., Giles, B. L., Wilder, F. D., et al. (2016). Magnetospheric Multiscale instrument suite operations and data system. *Space Sci. Rev.* 199, 545–575. doi: 10.1007/s11214-014-0128-5
- Bergen, K. J., Johnson, P. A., de Hoop, M. V., and Beroza, G. C. (2019). Machine learning for data-driven discovery in solid Earth geoscience. *Science* 363:eaau0323. doi: 10.1126/science.aau0323
- Berkson, J. (1944). Application of the logistic function to bio-assay. *J. Am. Stat. Assoc.* 39, 357–365. doi: 10.1080/01621459.1944.10500699
- Bobra, M. G., and Couvidat, S. (2015). Solar flare prediction using SDO/HMI vector magnetic field data with a machine-learning algorithm. *Astrophys. J.* 798:135. doi: 10.1088/0004-637X/798/2/135

## FUNDING

This material was based on work supported by the National Science Foundation Graduate Research Fellowship Program under Grant No. DGE 1256260 and was partially funded by the Michigan Space Grant Consortium under NNX15AJ20H. JL received funding through an NICHD training grant to the Population Studies Center at the University of Michigan (T32HD007339). ML was funded by NASA grant NNX16AQ04G.

## ACKNOWLEDGMENTS

We would like to thank Monica Bobra, Brian Swiger, Garrett Limon, Kristina Fedorenko, Dr. Nils Smit-Anseeuw, and Dr. Jacob Bortnik for relevant discussions related to this draft. We would also like to thank the conference organizers of the 2019 Machine Learning in Heliophysics conference at which this work was presented, and the American Astronomical Society Thomas Metcalf Travel Award for funding travel to this conference. This work has additionally appeared as a dissertation chapter (Azari, 2020). **Figure 2’s** copyright is held by Falconieri Visuals. It is altered here with permission. **Figure 1** contains graphics from Jia et al. (2012) and Chen et al. (2019) which can be found in journals with copyright held by AGU. We would like to thank Dr. Jon Vandegriff for assistance with the CHEMS data used within this work.

- Brand, A., Allen, L., Altman, M., Hlava, M., and Scott, J. (2015). Beyond authorship: attribution, contribution, collaboration, and credit. *Learn. Publish.* 28, 151–155. doi: 10.1087/20150211
- Breiman, L. (2001). Random forests. *Mach. Learn.* 45, 5–32. doi: 10.1023/A:1010933404324
- Butler, K. T., Davies, D. W., Cartwright, H., Isayev, O., and Walsh, A. (2018). Machine learning for molecular and materials science. *Nature* 559, 547–555. doi: 10.1038/s41586-018-0337-2
- Camporeale, E. (2019). The challenge of machine learning in space weather: nowcasting and forecasting. *Space Weath.* 17, 1166–1207. doi: 10.1029/2018SW002061
- Chen, Y., and Hill, T. W. (2008). Statistical analysis of injection/dispersion events in Saturn’s inner magnetosphere. *J. Geophys. Res.* 113:A07215. doi: 10.1029/2008JA013166
- Chen, Y., Hill, T. W., Rymer, A. M., and Wilson, R. J. (2010). Rate of radial transport of plasma in Saturn’s inner magnetosphere. *J. Geophys. Res.* 115:A10211. doi: 10.1029/2010JA015412
- Chen, Y., Manchester, W. B., Hero, A. O., Toth, G., DuFumier, B., Zhou, T., et al. (2019). Identifying solar flare precursors using time series of SDO/HMI images and SHARP parameters. *Space Weath.* 17, 1404–1426. doi: 10.1029/2019SW002214
- Couronné, R., Probst, P., and Boulesteix, A. L. (2018). Random forest versus logistic regression: a large-scale benchmark experiment. *BMC Bioinformatics* 19:270. doi: 10.1186/s12859-018-2264-5
- Coveney, P. V., Dougherty, E. R., and Highfield, R. R. (2016). Big data need big theory too. *Philos. Trans. R. Soc. A* 374. doi: 10.1098/rsta.2016.0153
- Cramer, J. S. (2002). *The Origins of Logistic Regression*. Tinbergen Institute Working Paper No. 2002-119/4.
- DeJong, A. D., Burch, J. L., Goldstein, J., Coates, A. J., and Young, D. T. (2010). Low-energy electrons in Saturn’s inner magnetosphere

- and their role in interchange injections. *J. Geophys. Res.* 115:A10229. doi: 10.1029/2010JA015510
- Dougherty, M. K., Kellock, S., Southwood, D. J., Balogh, A., Smith, E. J., Tsurutani, B. T., et al. (2004). The Cassini magnetic field investigation. *Space Sci. Rev.* 114, 331–383. doi: 10.1007/s11214-004-1432-2
- Ebert-Uphoff, I., Samarasinghe, S. M., and Barnes, E. A. (2019). Thoughtfully using artificial intelligence in Earth science. *Eos Transactions American Geophysical Union.* 100. doi: 10.1029/2019EO135235
- Fawcett, T. (2006). An introduction to roc analysis. *Pattern Recogn. Lett.* 27, 861–874. doi: 10.1016/j.patrec.2005.10.010
- Heidke, P. (1926). Berechnung des Erfolges und der Güte der Windstärkevorhersagen im Sturmwarnungsdienst (Calculation of the success and goodness of strong wind forecasts in the storm warning service). *Geografiska Annaler Stockholm* 8, 301–349. doi: 10.1080/20014422.1926.11881138
- Iten, R., Metger, T., Wilming, H., del Rio, L., and Renner, R. (2020). Discovering physical concepts with neural networks. *Phys. Rev. Lett.* 124:10508. doi: 10.1103/PhysRevLett.124.010508
- Jia, X., Hansen, K. C., Gombosi, T. I., Kivelson, M. G., Tóth, G., Dezeewu, D. L., et al. (2012). Magnetospheric configuration and dynamics of Saturn's magnetosphere: a global MHD simulation. *J. Geophys. Res.* 117:A05225. doi: 10.1029/2012JA017575
- Karpatne, A., Ebert-Uphoff, I., Ravela, S., Babaie, H. A., and Kumar, V. (2019). Machine learning for the geosciences: challenges and opportunities. *IEEE Trans. Knowl. Data Eng.* 31, 1544–1554. doi: 10.1109/TKDE.2018.2861006
- Kennelly, T. J., Leisner, J. S., Hospodarsky, G. B., and Gurnett, D. A. (2013). Ordering of injection events within Saturnian SLS longitude and local time. *J. Geophys. Res.* 118, 832–838. doi: 10.1002/jgra.50152
- Kokar, M. M. (1986). "Coper: A methodology for learning invariant functional descriptions," in *Machine Learning*. The Kluwer International Series in Engineering and Computer Science (Knowledge Representation, Learning and Expert Systems), eds T. M. Mitchell, J. G. Carbonell, and R. S. Michalski (Boston, MA: Springer), 151–154. doi: 10.1007/978-1-4613-2279-5\_34
- Krimigis, S. M., Mitchell, D. G., Hamilton, D. C., Livi, S., Dandouras, J., Jaskulek, S., et al. (2004). Magnetosphere Imaging Instrument (MIMI) on the Cassini mission to Saturn/Titan. *Space Sci. Rev.* 114, 233–329. doi: 10.1007/s11214-004-1410-8
- Lai, H. R., Russell, C. T., Jia, Y. D., Wei, H. Y., and Dougherty, M. K. (2016). Transport of magnetic flux and mass in Saturn's inner magnetosphere. *J. Geophys. Res.* 121, 3050–3057. doi: 10.1002/2016JA022436
- Liemohn, M. W., McCollough, J. P., Jordanova, V. K., Ngwira, C. M., Morley, S. K., Cid, C., et al. (2018). Model evaluation guidelines for geomagnetic index predictions. *Space Weath.* 16, 2079–2102. doi: 10.1029/2018SW02067
- Lockhart, J. W., and Weiss, G. M. (2014a). "Limitations with activity recognition methodology and data sets," in *Proceedings of the 2014 ACM International Joint Conference on Pervasive and Ubiquitous Computing: Adjust Publication* (Seattle, WA), 747. doi: 10.1145/2638728.2641306
- Lockhart, J. W., and Weiss, G. M. (2014b). "The benefits of personalized smartphone-based activity recognition models," in *2014 SIAM International Conference on Data Mining* (Philadelphia, PA). doi: 10.1137/1.9781611973440.71
- Manzato, A. (2005). An odds ratio parameterization for ROC diagram and skill score indices. *Weath. Forecast.* 20, 918–930. doi: 10.1175/WAF899.1
- Molnar, C. (2019). *Interpretable Machine Learning. A Guide for Making Black Box Models Explainable*. Available online at: <https://christophm.github.io/interpretable-ml-book/>
- Morley, S. K. (2020). Challenges and opportunities in magnetospheric space weather prediction. *Space Weath.* 18:e2018SW002108. doi: 10.1029/2018SW002108
- Müller, A. L., Saur, J., Krupp, N., Roussos, E., Mauk, B. H., Rymer, A. M., et al. (2010). Azimuthal plasma flow in the Kronian magnetosphere. *J. Geophys. Res.* 115:A08203. doi: 10.1029/2009JA015122
- NASA Headquarters (1980). *Voyager Backgrounder, Release No: 80-160*. NASA Headquarters. Available online at: <https://ntrs.nasa.gov/archive/nasa/casi.ntrs.nasa.gov/19810001583.pdf>
- NASA Jet Propulsion Laboratory (2017a). *Cassini Huygens by the Numbers*. NASA Jet Propulsion Laboratory. Available online at: <https://solarsystem.nasa.gov/resources/17761/cassini-huygens-by-the-numbers/>
- NASA Jet Propulsion Laboratory (2017b). *Mars Reconnaissance Orbiter By the Numbers*. NASA Jet Propulsion Laboratory. Available online at: <https://mars.nasa.gov/resources/7741/mars-reconnaissance-orbiter-by-the-numbers/>
- Pedregosa, F., Varoquaux, G., Gramfort, A., Michel, V., Thirion, B., Grisel, O., et al. (2011). Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830. Available online at: <https://dl.acm.org/doi/10.5555/1953048.2078195>
- Planetary Society (2020). *Missions to Mars*. Planetary Society. Available online at: <https://www.planetary.org/explore/space-topics/space-missions/missions-to-mars.html>
- Raissi, M., Perdikaris, P., and Karniadakis, G. E. (2019). Physics-informed neural networks: A deep learning framework for solving forward and inverse problems involving nonlinear partial differential equations. *J. Comput. Phys.* 378, 686–707. doi: 10.1016/j.jcp.2018.10.045
- Ren, H., Stewart, R., Song, J., Kuleshov, V., and Ermon, S. (2018). Learning with weak supervision from physics and data-driven constraints. *AI Mag.* 39, 27–38. doi: 10.1609/aimag.v39i1.2776
- Rudin, C. (2019). Stop explaining black box machine learning models for high stakes decisions and use interpretable models instead. *Nat. Mach. Intell.* 1, 206–215. doi: 10.1038/s42256-019-0048-x
- Swiger, B. M., Liemohn, M. W., and Ganushkina, N. Y. (2020). Improvement of plasma sheet neural network accuracy with inclusion of physical information. *Front. Astron. Space Sci.* doi: 10.3389/fspas.2020.00042
- Waskom, M., Botvinnik, O., Ostblom, J., Lukauskas, S., Hobson, P., Gelbart, M., et al. (2020). *mwaskom/seaborn: v0.10.0* (January 2020). doi: 10.5281/zenodo.3629446

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

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