

# Detection of Environment Transitions in Time Series Data for Responsive Science

Ameya Daigavane  
ameya.s.daigavane@jpl.nasa.gov  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA

Corey Cochrane  
corey.j.cochrane@jpl.nasa.gov  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA

Kiri L. Wagstaff  
kiri.l.wagstaff@jpl.nasa.gov  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA

Caitriona Jackman  
cjackman@cp.dias.ie  
Dublin Institute for Advanced Studies  
Dublin, Ireland

Gary Doran  
gary.b.doran.jr@jpl.nasa.gov  
Jet Propulsion Laboratory  
California Institute of Technology  
Pasadena, CA

Abigail Rymer  
abigail.rymer@jhuapl.edu  
Johns Hopkins University  
Applied Physics Laboratory  
Laurel, MD

## ABSTRACT

Onboard, *in-situ* detection of interesting events in time series data can enable the instrument or sensor to take action in a responsive way. Multivariate time series capture a richer depiction of the environment at each time step and therefore facilitate the detection of more subtle or complex events. In this work, we focus on the detection of events that mark the transition between different physical environments, as observed by a moving sensor such as a spacecraft. Our approach is to model such events as anomalies and employ anomaly detection and hidden state modeling methods to detect them. We assess RuLSIF, HOT SAX, HMM, and Matrix Profile methods and propose a novel extension of the Matrix Profile to multivariate data. We apply these methods to observations of the magnetosphere collected by the Cassini spacecraft in orbit around Saturn. Reliable detection of environment transitions could enable future instruments, such as the Plasma Instrument for Magnetic Sounding (PIMS) instrument on the Europa Clipper spacecraft, to adaptively select the best observing mode according to the current environment, yielding higher quality data.

## KEYWORDS

time series analysis, boundary detection, change detection, anomaly detection, CAPS, PIMS, Europa Clipper

### ACM Reference Format:

Ameya Daigavane, Kiri L. Wagstaff, Gary Doran, Corey Cochrane, Caitriona Jackman, and Abigail Rymer. 2020. Detection of Environment Transitions in Time Series Data for Responsive Science. In *MileTS '20: 6th KDD Workshop on Mining and Learning from Time Series, August 24th, 2020, San Diego, California, USA*. ACM, New York, NY, USA, 5 pages.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from [permissions@acm.org](mailto:permissions@acm.org).

*MileTS '20, August 24th, 2020, San Diego, California, USA*

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

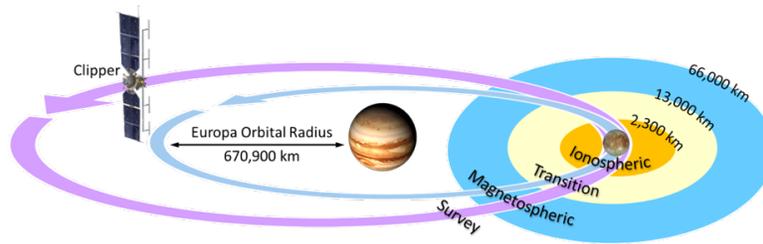
## 1 INTRODUCTION

Time series data consist of a sequence of measurements from a sensor that can be used to infer properties about the surrounding environment. For a sensor in motion, such as an instrument on a spacecraft, changes in the surrounding environment may manifest as detectable changes in the acquired observations. Online analysis of the time series data provides the basis for an autonomous system to decide when to take action due to a change in environment.

We focus on methods for the onboard detection of transition events that signal when the sensor has moved into a different operating environment. Many instruments have configurable operating modes, each of which is best suited for a particular environment. Detecting when the environment has changed allows the instrument to switch modes and optimize the quality of the collected data.

A concrete example of a spacecraft instrument that could benefit from onboard detection of environment transition events is the Plasma Instrument for Magnetic Sounding (PIMS) instrument on the upcoming Europa Clipper mission [4, 12]. PIMS will characterize the plasma around Jupiter's moon Europa to understand the influence of Europa's ocean on its magnetosphere by recording the distribution of energies of ions and other charged particles in the plasma. PIMS will have four observing modes (survey, magnetospheric, transition, and ionospheric) that are optimized for the anticipated plasma conditions in each region (Figure 1). Since Europa is too far away from Earth to allow realtime human control, mission planners will pre-specify the instrument's operating mode at each part of the orbit in advance. Uncertainty in the prior knowledge of the location of these boundaries has led to the planned use of a "transition" mode that alternates between observing modes as a compromise: it is ideal for neither one but ensures that at least some useful data is collected. Onboard, *in situ* detection of the transitions could enable the instrument to switch adaptively to the best observing mode and reduce the need for the transition mode. The result would likely be the opportunity to capture a greater number of electromagnetic phenomena of scientific interest, at the best range of observing energies.

We propose the use of anomaly detection and state transition methods to automatically detect environment transitions. We apply



**Figure 1: PIMS will operate in magnetospheric mode far from Europa, ionospheric mode close to Europa, and transition mode in between. Note that this figure is not drawn to scale.**

these methods to data collected by the Cassini Plasma Spectrometer (CAPS) Electron Spectrometer (ELS) instrument at Saturn as an analogue for the data to be collected by PIMS at Europa (Section 2). In Section 3, we describe the algorithms that we expect to be suitable for use in the computationally limited onboard environment. We report empirical results on the task of detecting magnetospheric crossings in CAPS-ELS data (Section 5) and in Section 6 we describe the next steps to realizing the goal of onboard deployment of the best performing methods.

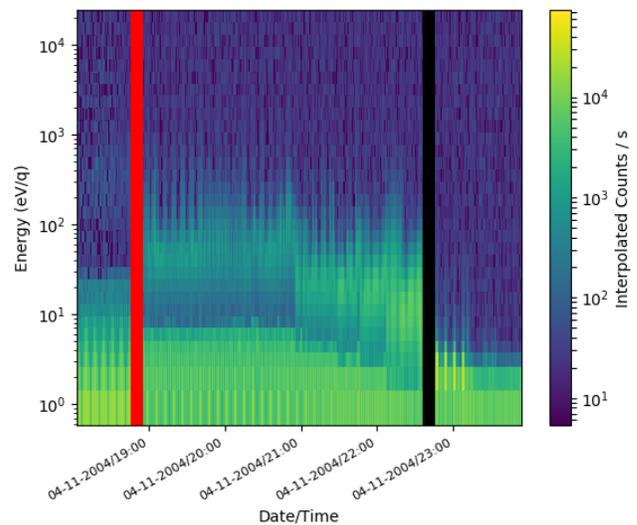
## 2 TIME SERIES DATA FROM MAGNETOSPHERIC INSTRUMENTS

PIMS is currently under development, and data from the instrument does not yet exist. For this study, we employed data from CAPS-ELS, an analogous instrument on the recently completed Cassini mission. CAPS-ELS measured electron flux as a function of their energy and direction [16]. The multivariate time series contains observations roughly every 2 seconds that consist of electron counts tallied in energy bins that range from 0.53 to 28200 eV.

Data from CAPS-ELS spans the years 2004 to 2012 and includes 2399 known (manually annotated) magnetic field boundary crossings [5]<sup>1</sup>. These boundaries correspond to Saturn’s magnetopause (between the magnetosphere and magnetosheath) and bow shock (between the magnetosheath and the solar wind). Although these boundaries are unlikely to have exact analogues in the near-Europa environment, the principles developed here will help identify boundaries that Europa Clipper will observe with PIMS.

### 2.1 Pre-processing ELS Data

We employed several pre-processing steps to prepare CAPS-ELS data for analysis. These pre-processing steps would not be needed if we had access to the original data in the onboard setting. Following the recommendations of the CAPS PDS User Guide [13], we use data from the least-obstructed anode 5. The original CAPS-ELS data consists of observations at 63 energy bins, but some observations were reduced to 32 bins prior to downlink to reduce bandwidth consumption. For consistent dimensionality, we excluded the lowest-energy bin entirely (due to residual spurious signals) and converted all remaining 62-bin observations to 31-bin observations by storing the average value for pairs of adjacent bins.



**Figure 2: Example 6-hour CAPS-ELS observation containing both bow shock (black) and magnetopause (red) crossings.**

CAPS-ELS data also has an irregular sampling cadence because some observations adjacent in time were combined prior to downlink. Further, around 0.6% of individual bin observations are missing due to downlink gaps. To address both issues, we performed a linear spline interpolation across the time dimension independently for each of the 31 bins, at each equally spaced time step. Any negative counts obtained after interpolation were clipped to 0. On our training set spanning 180 hours of observations from 2004, fewer than 0.1% of the entries were clipped. The resulting 31-dimensional data set has a regular sampling cadence of 2 seconds.

To smooth the time series, we convolved it with a Gaussian kernel, removing high-frequency components. To remove additive noise, we used either a minimum, median, or maximum filter. The width of the Gaussian kernel and the filter type were chosen per algorithm via validation. Finally, we employed Anscombe’s transform to each binned count  $c$  to convert the observations (count-based Poisson distribution) to a variable with an approximately Gaussian distribution,  $c \rightarrow \sqrt{c + \frac{1}{3}}$ .

<sup>1</sup>Labeled data set: <https://doi.org/10.5281/zenodo.3946033>

## 2.2 Events: Magnetic Field Boundary Crossings

As the Cassini spacecraft orbited Saturn, it repeatedly crossed into and out of the magnetosphere, magnetosheath, and solar wind regions. There are two event types that we seek to detect in the CAPS-ELS data: the magnetopause (transition between magnetosphere and magnetosheath) and the bow shock (between magnetosheath and solar wind). As illustrated in Figure 2, bow shock crossings are often more prominent and easier to detect.

Each of these crossings can be further broken down into two categories, IN and OUT, based how the spacecraft moved relative to the boundary. The numbers of events of each type observed during Cassini’s mission, broken down by year, are as follows:

Year	BS IN	BS OUT	MP IN	MP OUT	Total
2004	18	17	13	12	60
2005	98	97	102	101	398
2006	6	6	49	49	110
2007	91	92	254	253	690
2008	16	16	128	126	286
2009	2	2	86	86	176
2010	23	23	142	142	330
2011	45	45	60	60	210
2012	20	19	21	19	79
All	319	317	855	848	2339

## 3 TRANSITION DETECTION ALGORITHMS

Our approach is to compute a transition detection score for each observation in the time series. The baseline algorithm  $B$  computes the Euclidean distance between observation  $t_i \in \mathbb{R}^d$  at time  $i$  and the immediately preceding observation  $t_{i-1}$ :

$$B_i = \|t_i - t_{i-1}\|_2 \quad (1)$$

We assessed four unsupervised approaches to detecting transitions either by modeling states, identifying change points, or detecting anomalies. A key constraint for *in-situ* data analysis is the limited computational resources available on-board the spacecraft, which restricts our options to algorithms that have low memory and computational requirements.

First, we used a Hidden Markov Model (HMM)<sup>2</sup> to identify changes in state. Each state  $s_j$  is modeled by a  $d$ -dimensional Gaussian distribution with mean  $\mu_{s_j}$  and covariance matrix  $\Sigma_{s_j}$ , learned from the data. We compute a dissimilarity matrix  $DS \in \mathbb{R}^{s \times s}$ , where  $s$  is the number of states and each entry is defined as:

$$DS_{s_1, s_2} = KL(\mathcal{N}(\mu_{s_1}, \Sigma_{s_1}) \parallel \mathcal{N}(\mu_{s_2}, \Sigma_{s_2})) + KL(\mathcal{N}(\mu_{s_2}, \Sigma_{s_2}) \parallel \mathcal{N}(\mu_{s_1}, \Sigma_{s_1}))$$

using the symmetrized  $KL$ -divergence between the learned distributions. The HMM anomaly detection score at time step  $i$  is

$$HMM_i = \delta_i^T DS \delta_i \quad (2)$$

where  $\delta_i = |p_i - p_{i-1}|$ , the absolute value of the difference in posterior state distributions  $p_i \in \mathbb{R}^s$ . The use of the dissimilarity matrix  $DS$  assigns more weight to transitions between dissimilar states. We use both Bayesian and non-Bayesian HMM models: the parameters for the non-Bayesian (“vanilla”) HMM are learned via Expectation Maximization [1], while the Bayesian HMM models

<sup>2</sup>See <https://github.com/hmmlearn/hmmlearn> and <https://github.com/mattjj/pyhsmm/>

the joint distribution over input samples and latent variables as a “sticky” Hierarchical Dirichlet Process [3] and then averages over 1000 samples of latent variables obtained via Gibbs sampling [11].

We also employed RuLSIF<sup>3</sup> (Relative Unconstrained Least Squares Information Fitting [7] which seeks to identify “change points” in a time series. The RuLSIF score  $RS$  is the symmetrized  $\alpha$ -relative PE-divergence between the distribution  $P_i$  of  $m$  observations after the current time-step  $i$  and the distribution  $P'_i$  of  $m$  observations before time-step  $i$ .

$$RS_i = PE_\alpha(P_i \parallel P'_i) + PE_\alpha(P'_i \parallel P_i) \quad (3)$$

As in [7], we “pack” the time series by taking subsequences of length  $k$  rather than a single time step at a time. This increases the dimensionality of each sample by a factor of  $k$ .

Finally, we employed two methods for discord or anomaly detection in time series. The first is HOT SAX<sup>4</sup> (Heuristically Ordered Time series using Symbolic Aggregate Approximation) [6]. Let  $X \in \mathbb{R}^{n \times d}$  be a  $d$ -dimensional time series of length  $n$  and  $X_{i,w} \in \mathbb{R}^{w \times d}$  be a  $z$ -normalized subsequence of length  $w$  starting at time  $i$ . Then the HOT SAX score is the minimum distance between the current subsequence and every other non-overlapping subsequence of the same length.

$$HS_i = \min_{\substack{|i-j| \geq w \\ 0 \leq j \leq n-w}} \|X_{i,w} - X_{j,w}\|_F, \quad (4)$$

where  $\|\cdot\|_F$  represents the Frobenius norm. To compute SAX representations of multidimensional subsequences, we use SAX-ZSCORE [9]. To speed up HOT SAX, we reduced the dimensionality of the time series using Principal Component Analysis (PCA) to reduce the data to a lower-dimensional subspace while capturing as much of the variance as possible.

The second anomaly detection method is based on the Matrix Profile [15] with a novel extension to operate on multidimensional data and maximize sensitivity to anomalies<sup>5</sup>. For a unidimensional time series  $T$  of length  $n$ , let  $T_{i,w} \in \mathbb{R}^w$  be the subsequence of length  $w$  starting at time  $i$ . The Matrix Profile  $MP \in \mathbb{R}^{n-w+1}$  records, for each time  $i$ , the minimum Euclidean distance from the  $z$ -normalized subsequence  $T_{i,w}$  to all other non-overlapping  $z$ -normalized subsequences of length  $w$  in the time series [15].

$$MP_i = \min_{\substack{|i-j| \geq w \\ 0 \leq j \leq n-w}} \|T_{i,w} - T_{j,w}\|_2 \quad (5)$$

For a  $d$ -dimensional time series, we propose an anomaly-sensitive Multidimensional Matrix Profile  $MMP \in \mathbb{R}^{n-w+1}$  that consists of the sum across dimensions  $d$  of unidimensional Matrix Profiles  $MP^{(d)}$  computed independently for each dimension:

$$MMP_i = \sum_d MP_i^{(d)} \quad (6)$$

Subsequences with large nearest-neighbor distances in multiple dimensions will receive large scores. It is also possible to restrict sensitivity to the  $k$  largest matrix profile distances per time step. This approach contrasts with the  $k$ -dimensional matrix profile

<sup>3</sup>[https://github.com/hoxo-m/densratio\\_py](https://github.com/hoxo-m/densratio_py)

<sup>4</sup><https://github.com/ameya98/saxpy>

<sup>5</sup><https://github.com/JPLMLIA/CAPS-ELS-Transition-Detection>

**Table 1: Parameters for bow shock crossing detectors**

Algorithm	Parameters (No blur if not stated)
Baseline	Max Filter (Size 120), Center 10 bins.
HOT SAX	Window Size $w = 50$ , Discords = 2, PCA-5, Max Filter (Size 120), All 31 bins.
Matrix Profile	Window Size $w = 100$ , Noise Correction = 0.6, PCA-10, Median Filter (Size 120), Center 10 bins.
RuLSIF	Window Size = 20, Packing Factor $k = 1$ , Median Filter (Size 30), All 31 bins.
Bayesian HDP-HMM	Maximum States = 2, Sticky Coefficient = 0.1, PCA-5, Median Filter (Size 120), Center 10 bins, Gaussian Blur $\sigma = 5$ .
Non-Bayesian HMM	Number of States = 2, PCA-10, Median Filter (Size 120), All 31 bins, Gaussian Blur $\sigma = 5$ .

**Table 2: Parameters for magnetopause crossing detectors**

Algorithm	Parameters (No blur if not stated)
Baseline	Max Filter (Size 120), Center 10 bins, Gaussian Blur $\sigma = 5$ .
HOT SAX	Window Size $w = 400$ , Discords = 2, PCA-5, Median Filter (Size 120), All 31 bins.
Matrix Profile	Window Size $w = 100$ , Noise Correction = 0.2, PCA-10, Max Filter (Size 120), Center 10 bins.
RuLSIF	Window Size $w = 60$ , Packing Factor $k = 5$ , Median Filter (Size 120), All 31 bins.
Bayesian HDP-HMM	Maximum States = 3, Sticky Coefficient = 0.1, PCA-5, Median Filter (Size 120), Center 10 bins, Gaussian Blur $\sigma = 5$ .
Non-Bayesian HMM	Number of States = 2, PCA-5, Median Filter (Size 120), All 31 bins.

defined by Yeh et al. [14] which uses the  $k$  smallest values to emphasize similarity. In this study, we employed all dimensions. We also incorporated noise correction [10], an improvement that subtracts out the effect of Gaussian noise in the Matrix Profile computation, without changing its runtime asymptotically.

This algorithm can be applied to any multidimensional time series. Computing the Matrix Profiles for each dimension takes  $O(n^2)$  time per dimension, for a total of  $O(n^2d)$  time, and computing the sum across dimensions at each time step takes  $O(nd)$  time.

#### 4 METHODOLOGY AND METRICS

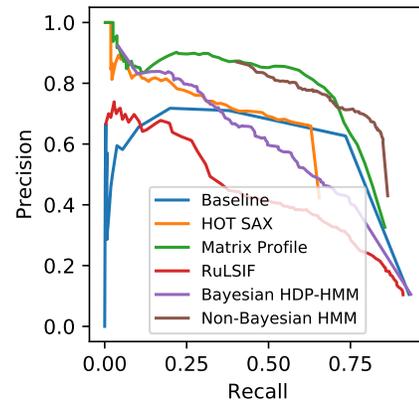
For each event type, we selected the best parameters for each algorithm (see Table 1 and Table 2) using a training set of all observations from 2004, which encompass 60 crossing events. Observations from 2005 through 2012 comprise our held-out test set.

We conducted a retrospective evaluation by measuring the reliability of the detection scores generated by each algorithm at each time step. We assessed recall and precision by defining a time tolerance of  $t_{tol}$  minutes and, for each algorithm, generated a list of detections that exceeded a given score threshold  $s_{thres}$  and then retained only local peak detections (those with scores greater than all detections within  $\pm \frac{t_{tol}}{2}$  seconds). We defined true positives as detections within a window that spans 1 minute before and  $t_{tol}$  minutes after a labeled event, false positives as detections outside

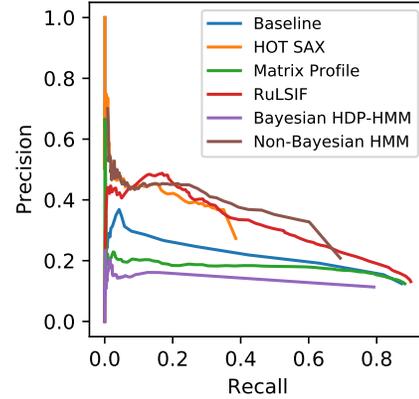
of all such windows, and false negatives as labeled events without any detections within the sensitivity window.

To assess feasibility, we also computed the average runtime for each algorithm. RuLSIF was the most computationally expensive method by far, requiring an average of 40 minutes on a 2.1-GHz processor to process a single 6-hour observation for bow shock detection and more than 5 hours for magnetopause detection (due to larger window size  $w$  and packing factor  $k$  parameters). For onboard operation we would likely restrict these parameter options. In contrast, HOT SAX and both HMM methods required less than a minute, and the Matrix Profile took an average of about 8 minutes.

#### 5 EXPERIMENTAL RESULTS



(a) Bow shock



(b) Magnetopause

**Figure 3: Test performance on data from 2005 to 2012**

Results for detecting events in data from 2005 through 2012 using a  $t_{tol}$  value of 20 minutes are shown in Figure 3. We found that the bow shock crossings were much easier than magnetopause crossings to reliably detect, as shown in Figure 3(a). In this case, the best-performing method was the Matrix Profile, with the non-Bayesian HMM as the next-best performing method. For this

application, in which a spacecraft instrument will change its observing mode based on event detections, precision is very important since false detections will cause unwanted changes. At a required precision of 0.8, the Matrix Profile achieved a recall of 0.636, and the non-Bayesian HMM had a recall of 0.583.

However, performance varies significantly across different years of the mission, indicating that concept drift is a challenge for generalization. For 2012, both the Matrix Profile and the non-Bayesian HMM achieved a recall of 0.923, while in 2006, HOT SAX achieved a recall of 0.750 yet the Matrix Profile (0.083) and non-Bayesian HMM (0.583) recall scores were much lower (given a precision threshold of 0.8). We are investigating the use of online methods to update parameter settings based on the most recent observations to improve generalization.

For magnetopause detection, the overall best-performing method was the non-Bayesian HMM. Unfortunately, no methods achieved 0.8 precision across all years (Figure 3(b)). The best result observed was for 2006, in which HOT SAX achieved a recall of 0.146 (at precision 0.8). MP events are characterized by more diffuse boundaries and can be difficult to identify visually from CAPS-ELS data alone. The labeled events that we employed were identified originally from magnetometer data which can capture those events more clearly. We are also looking into whether supervised methods that employ the labeled data directly, rather than approaching the problem as an anomaly detection task, can improve performance.

## 6 CONCLUSIONS AND FUTURE WORK

We have demonstrated the detection of magnetic field boundaries using data from spectrometers such as CAPS-ELS. To our knowledge, this is the first work of its kind. Next, we will develop prototypes of these algorithms in the C programming language and test them using a PPC750 processor, which has capabilities similar to those of the processor planned for use on the Europa Clipper spacecraft. We will employ the BITFLIPS software radiation simulator to assess radiation sensitivity using the methodology employed for evaluating thermal anomaly detectors in the harsh European radiation environment [2]. We will also assess algorithm resource consumption given the spacecraft's existing activities and available resources using the APGen [8] activity plan generator and simulation tool.

We are also interested in assessing time series detection in an online scenario. When making onboard decisions about when to change an instrument's observing mode, detections must be made with respect to only the previously collected observations rather than the full time series, as employed in the retrospective study reported in this work. The online version of the detectors would likely employ a different peak-finding strategy, and it may also be beneficial to create different detectors for transitions IN versus OUT of a given environment.

To date our experiments have been done with analogue data collected by the CAPS-ELS instrument at Saturn, but we do not expect identical empirical results for data collected by PIMS at Europa. For operational use, we would plan to collect PIMS data from one or more initial flybys of Europa to enable an offline algorithm evaluation customized to the boundary events encountered by the Europa Clipper spacecraft. The same methodology employed in the

CAPS-ELS study can be used to identify the best-performing algorithm on PIMS Europa data, in conjunction with our assessment of the runtime cost and resource consumption.

## ACKNOWLEDGMENTS

The High Performance Computing resources used in this investigation were provided by funding from the JPL Office of the Chief Information Officer. Part of this research was carried out at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration. Copyright 2020. All rights reserved.

## REFERENCES

- [1] Leonard E. Baum and Ted Petrie. 1966. Statistical inference for probabilistic functions of Finite State Markov Chains. *The Annals of Mathematical Statistics* 37, 6 (1966), 1554–1563.
- [2] Gary Doran, Kiri L. Wagstaff, Jonathan Bapst, Ashley G. Davies, and Saadat Anwar. 2020. Assessing the susceptibility of Europa Clipper onboard thermal anomaly detection to radiation. In *51st Lunar and Planetary Science Conference*. Lunar and Planetary Institute, Abstract 2193.
- [3] E. B. Fox, E. B. Sudderth, M. I. Jordan, and A. S. Willsky. 2008. An HDP-HMM for Systems with State Persistence. In *Proceedings of the 25th International Conference on Machine Learning (ICML)*. 312–319.
- [4] Matthew Grey, Joseph Westlake, Shawn Liang, Erik Hohlfeld, Alexander Crew, and Ralph McNutt. 2018. Europa PIMS prototype Faraday cup development. In *Proceedings of the 2018 IEEE Aerospace Conference*. IEEE. <https://doi.org/10.1109/AERO.2018.8396522>
- [5] Caitriona M. Jackman, Michele K. Dougherty, and M. F. Thomsen. 2019. Survey of Saturn's magnetopause and bow shock positions over the entire Cassini mission: boundary statistical properties, and exploration of associated upstream conditions. *Journal of Geophysical Research: Space Physics* 124, 11 (2019), 8865–8883. <https://doi.org/10.1029/2019JA026628>
- [6] Eamonn Keogh, Jessica Lin, and Ada Fu. 2005. HOT-SAX: Efficiently finding the most unusual time series subsequence. In *Fifth IEEE International Conference on Data Mining (ICDM'05)*. IEEE, 226–233.
- [7] Song Liu, Makoto Yamada, Nigel Collier, and Masashi Sugiyama. 2013. Change-point detection in time-series data by relative density-ratio estimation. *Neural Networks* 43 (2013), 72–83. <https://doi.org/10.1016/j.neunet.2013.01.012>
- [8] Pierre F. Mardague, Steven S. Wissler, Matthew D. Lenda, and Daniel F. Finnerty. 2014. APGEN scheduling: 15 years of experience in planning automation. In *SpaceOps 2014 Conference*. <https://doi.org/10.2514/6.2014-1809>
- [9] Y. Mohammad and T. Nishida. 2014. Robust Learning from Demonstrations Using Multidimensional SAX. In *14th International Conference on Control, Automation and Systems (ICCAS 2014)*. 64–71.
- [10] D. De Paepe, O. Janssens, and S. Van Hoecke. 2019. Eliminating Noise in the Matrix Profile. In *Proceedings of the 8th International Conference on Pattern Recognition Applications and Methods*. 84–93.
- [11] Yee Whye Teh, Michael I. Jordan, Matthew J. Beal, and David M. Blei. 2006. Hierarchical Dirichlet Processes. *J. Amer. Statist. Assoc.* 101, 476 (2006), 1566–1581.
- [12] J. H. Westlake, R. L. McNutt, J. C. Kasper, A. W. Case, M. P. Grey, C. K. Kim, C. C. Battista, A. Rymer, C. S. Paty, X. Jia, M. L. Stevens, K. Khurana, M. G. Kivelson, J. A. Slavin, H. H. Korth, H. T. Smith, N. Krupp, E. Roussos, and J. Saur. 2016. The Plasma Instrument for Magnetic Sounding (PIMS) on the Europa Clipper mission. In *AAS/Division for Planetary Sciences Meeting Abstracts #48*, Vol. 48. 123.27.
- [13] R. J. Wilson, F. Cray, L. K. Gilbert, D. B. Reisenfeld, J. T. Steinberg, and R. Livi. 2012. Cassini Plasma Spectrometer (CAPS) PDS User's Guide. [https://pds-atmospheres.nmsu.edu/data\\_and\\_services/atmospheres\\_data/Cassini/inst-caps.html](https://pds-atmospheres.nmsu.edu/data_and_services/atmospheres_data/Cassini/inst-caps.html)
- [14] Chin-Chia Michael Yeh, Nickolas Kavantzias, and Eamonn J. Keogh. 2017. Matrix Profile VI: Meaningful multidimensional motif discovery. In *IEEE International Conference on Data Mining (ICDM)*. IEEE, 565–574.
- [15] Chin-Chia M. Yeh, Yan Zhu, Liudmila Ulanova, Nurjahan Begum, Yifei Ding, Hoang A. Dau, Diego F. Silva, Abdullah Mueen, and Eamonn Keogh. 2016. Matrix Profile I: All pairs similarity joins for time series: A unifying view that includes motifs, discords and shapelets. In *IEEE International Conference on Data Mining (ICDM)*. IEEE, 1317–1322.
- [16] D. T. Young, B. L. Barraclough, J. J. Berthelmer, M. Blanc, J. L. Burch, A. J. Coates, R. Goldstein, M. Grande, T. W. Hill, J. M. Illiano, M. A. Johnson, R. E. Johnson, R. A. Baragiola, V. Kelha, D. Linder, D. J. McComas, B. T. Narheim, J. E. Nordholt, A. Preece, E. C. Sittler, K. R. Svenes, S. Szalai, K. Szego, P. Tanskanen, and K. Viherkanto. 2004. Cassini Plasma Spectrometer investigation. *Space Science Reviews* 114, 1 (September 2004), 1–112.