

Article

Disentangling Spatial, Temporal, and Space Weather Contributions to Spacecraft Anomaly States: A Case–Control Analysis of GOES-16/17

Zongliang Li ^{1,†} , Tianyou Yu ^{2,†} , Danhuai Guo ^{2,*} , Fenglin Ding ¹, Yizhuo Liu ¹ and Xunchun Li ³

¹ Beijing Institute of Control Engineering, Beijing 100094, China; zongnet@126.com (Z.L.); cast_ding@163.com (F.D.); yizhuo.98@163.com (Y.L.)

² College of Information Science and Technology, Beijing University of Chemical Technology, Beijing 100029, China; yutianyou@buct.edu.cn

³ Academy of Broadcasting Science, National Radio and Television Administration, Beijing 100866, China; lixunchun@abs.ac.cn

* Correspondence: gdh@buct.edu.cn

† These authors contributed equally to this work.

Abstract

Recurring spacecraft anomalies may reflect anomaly-prone system states shaped jointly by spacecraft geometry, recent anomaly history, and space weather forcing, rather than isolated responses to single external drivers. Distinguishing these contributions is important for interpretable anomaly monitoring in geostationary orbit. This study develops a satellite-stratified case–control framework for GOES-16/17 Extreme Ultraviolet and X-ray Irradiance Sensors (EXIS) Space Wire (SpW) anomaly records. Orbital–illumination descriptors, same-satellite event history variables, and space weather variables from the NASA OMNI database are integrated within chronological train–test validation, supported by null-baseline comparison, stratified bootstrap confidence intervals, exclusion window sensitivity analysis, feature group ablation, cross-satellite testing, and calibration diagnostics. Orbital–illumination variables provide weak but reproducible discrimination, event history descriptors capture temporal clustering, and space weather variables add complementary held-out information. The full Space Environment-Integrated Model (SEIM) reached a test area under the receiver operating characteristic curve (AUC) of 0.7540, while a compact train-only L1-selected clean Top-15 model achieved comparable held-out discrimination with a test AUC of 0.7588 after excluding direct near-neighbor history variables. Bootstrap comparisons indicate that this small difference is not statistically significant. Calibration diagnostics further confirm that fitted scores should be interpreted as discriminative anomaly-state indicators under the case–control design rather than as calibrated operational anomaly probabilities.

Keywords: spacecraft anomaly; GOES; SpaceWire; space weather; case–control design; anomaly-state characterization; logistic regression; model validation; orbital–illumination geometry; event history descriptors



Academic Editor: Andrea Lani

Received: 25 May 2026

Revised: 19 June 2026

Accepted: 25 June 2026

Published: 27 June 2026

Copyright: © 2026 by the authors.

Licensee MDPI, Basel, Switzerland.

This article is an open access article distributed under the terms and conditions of the [Creative Commons Attribution \(CC BY\)](https://creativecommons.org/licenses/by/4.0/) license.

1. Introduction

A central objective of spacecraft prognostics and health management (PHM) is not only to diagnose anomalies after they occur, but also to recognize when a spacecraft or subsystem is entering an anomaly-prone state. For long-duration operational missions, recurring anomalies rarely occur as purely independent engineering events. Their timing may reflect

the combined influence of external environmental forcing, orbital and illumination geometry, and the recent operational or anomaly history of the system. A PHM-oriented anomaly analysis should therefore move beyond retrospective event association and ask a more general question: whether the observed anomaly sequence contains a temporally generalizable anomaly-state representation that can support interpretable anomaly-state monitoring.

This question is especially challenging for geostationary spacecraft. At geostationary orbit, the local spacecraft environment is shaped by both space weather forcing and orbital context. An apparent local-time dependence in anomaly occurrence may arise from the spacecraft's position relative to magnetospheric structure, from illumination or shadow-related conditions, or from local-time-dependent plasma and energetic particle populations. Similarly, a burst or cluster of anomalies may indicate persistent external disturbance, charging accumulation and relaxation, subsystem recovery behavior, or a broader susceptible state of the spacecraft–environment system. These mechanisms can produce similar timing patterns in the anomaly record. Without explicit separation, the same anomaly sequence may be interpreted as a geometric effect, a temporal clustering effect, or a space weather-driven effect, even though all of these components may contribute simultaneously.

Previous studies have established that geostationary spacecraft anomalies are associated with energetic particle conditions, geomagnetic activity, and local-time-dependent environments [1–3]. Charging-focused studies further show that anomaly risk can increase under particular plasma, solar wind, and geomagnetic conditions, especially when charging-relevant environments persist over finite timescales [4–7]. These results provide strong evidence that spacecraft anomaly occurrence is structured rather than random. However, for PHM applications, the key issue is not only whether anomalies are statistically associated with disturbed environments. It is also necessary to determine how much of the timing structure is attributable to orbital–illumination modulation, how much reflects recent anomaly-state persistence, and how much is attributable to external space weather forcing.

Recent anomaly surveys and reviews have emphasized that spacecraft anomaly analysis should move beyond highly aggregated catalogs and should consider orbit regime, anomaly type, subsystem context, and physical mechanism [8,9]. This direction is important because subsystem-level anomaly modes are closer to the actual units of monitoring, diagnosis, and operational response than mission-wide anomaly counts. Broad anomaly catalogs are valuable for identifying general space weather associations, but they may obscure whether a specific recurring anomaly class has a stable and temporally generalizable timing structure. For PHM, the more useful representation is a compact and physically interpretable description of the anomaly-prone state, rather than a purely retrospective correlation between events and environmental indices.

The GOES-R series provides a suitable validation setting for this problem because it combines long-duration geostationary operations, environmental monitoring, and well-documented instrument systems [10]. The Extreme Ultraviolet and X-ray Irradiance Sensors (EXIS) onboard GOES-R provide key solar X-ray and ultraviolet irradiance measurements for operational space weather monitoring [11]. The EXIS Space Wire anomalies observed on GOES-16 and GOES-17 form a focused recurring subsystem-level anomaly class in an operational geostationary context. Most directly, Kress et al. [12] showed that GOES-R operational anomalies, including EXIS Space Wire anomalies, are associated with enhanced electron fluxes and display recognizable local-time structure. This finding motivates a more general PHM question: whether the timing structure of this recurring anomaly class can be decomposed into spatial modulation, temporal clustering, and external forcing, and whether that decomposition generalizes to a later mission period.

In this study, we treat GOES-16/17 EXIS Space Wire anomalies as a validation case for PHM-oriented anomaly-state characterization. The analysis is built around three

complementary sources of timing information. The first is orbital–illumination structure, including local time, seasonal geometry, and shadow-related conditions. The second is same-satellite event history structure, which describes whether the spacecraft has recently been in an anomaly-active state. The third is external space environment forcing, including geomagnetic activity, solar wind conditions, interplanetary magnetic field structure, solar activity background, and particle environment proxies. These sources are not assumed to be mutually exclusive. Instead, they are treated as coupled components of the spacecraft–environment state: space weather forcing can persist over hours to days, local-time and illumination geometry can modulate the charging environment, and recent anomalies can provide an observable summary of continuing system susceptibility.

The objective of this paper is to disentangle these sources of anomaly timing structure within a common event-level framework and to evaluate whether the resulting information is temporally generalizable to later mission periods. We construct a satellite-stratified case–control data set from public GOES-16/17 EXIS Space Wire anomaly timestamps, ephemeris-derived orbital and illumination quantities, same-satellite event history descriptors, and OMNI space weather variables. Models are trained on an earlier portion of the record and evaluated on later samples using chronological train–test validation, with additional cross-satellite generalization tests. In this way, the reported discrimination reflects time-held-out generalization rather than apparent fit to the same mission period. We further assess robustness through exclusion window sensitivity, no-history ablation, drop-column and permutation-based group importance analyses, bootstrap uncertainty quantification, and calibration diagnostics.

1.1. Problem Statement

Given a set of recurring anomaly timestamps from two geostationary satellites (GOES-16 and GOES-17), the problem is to construct a discriminative anomaly-state representation that separates anomaly cases from temporally matched non-anomaly controls using physically interpretable geometry, event history, and space weather descriptors. The representation must generalize chronologically: models trained on earlier mission data should maintain discrimination on held-out later periods that were not used for fitting, preprocessing, or feature selection. The representation should also be evaluated for cross-satellite generalization and for sensitivity to the control sampling design. The target is anomaly-state discrimination and characterization within a case–control framework, not causal proof or calibrated operational probability estimation. Figure 1 provides a schematic overview of the proposed framework.




Figure 1. Problem statement and motivation of the proposed anomaly-state characterization framework. GOES-16/17 EXIS Space Wire anomaly records serve as event anchors. Similar anomaly timing

patterns may arise from orbital–illumination geometry, recent anomaly history, and space weather forcing. Therefore, the objective is not to attribute anomalies to a single factor, but to separate the relative contributions of geometry, event history, and space weather within a unified anomaly-state representation. The resulting framework aims to provide interpretable anomaly-state discrimination and to evaluate chronological and cross-satellite generalization, without estimating calibrated operational anomaly probabilities or establishing causality.

1.2. Work Contributions and Structure

The main contributions of this study are fourfold:

1. We formulate recurring spacecraft anomalies as a PHM-oriented anomaly-state characterization problem, emphasizing the need to separate orbital–illumination modulation, temporal clustering, and external space weather forcing within a subsystem-level anomaly process.
2. We develop a unified satellite-stratified case–control framework that time-aligns GOES-16/17 anomaly records, ephemeris-derived orbital–illumination geometry, same-satellite event history descriptors, and OMNI space weather variables, and evaluates temporal generalization through chronological train–test validation.
3. We systematically evaluate discrimination using chronological holdout, stratified bootstrap 95% confidence intervals, paired bootstrap model comparisons, a constant null baseline, exclusion window sensitivity analysis (± 0.5 h to ± 6 h), no-history ablation models, drop-column and permutation-based group importance diagnostics, cross-satellite generalization tests, lunar-geometric feature accounting, and calibration diagnostics with explicit caveats.
4. We derive a compact and physically interpretable selected-feature representation for anomaly-state monitoring, and explicitly distinguish predictive association under the case–control design from causality and calibrated operational probability.

The remainder of this paper is organized as follows. Section 2 reviews the related literature on geostationary spacecraft anomalies, charging environments, and PHM-oriented analysis. Section 3 describes the proposed method, including the case–control design, feature groups, baseline models, bootstrap uncertainty estimation, sensitivity and ablation analyses, cross-satellite tests, and calibration diagnostics. Section 4 presents the experimental results. Section 5 discusses the implications, interpretation boundaries, and limitations. Section 6 concludes the paper with advantages, limitations, and future perspectives.

2. Related Literature

Research on geostationary spacecraft anomalies has long shown that anomaly occurrence is not temporally uniform, but is organized by the interaction between spacecraft systems and the near-Earth space environment. Early statistical studies reported associations between GEO anomalies, geomagnetic activity, energetic electrons, and orbital context [1,2]. Event-based analyses further showed that anomaly rates vary with local time, storm phase, and particle conditions, indicating that anomaly timing reflects both external environmental forcing and spacecraft location within the geospace system [3]. Operationally oriented approaches, such as the SEAES-GEO framework, also emphasized that anomaly risk can be represented more effectively through physically meaningful environmental hazard indicators than through raw measurements alone [13]. These studies provide the foundation for treating GEO anomaly timing as a structured process rather than a random sequence of isolated events.

A major physical pathway linking the space environment to GEO anomalies is spacecraft charging. Reviews and empirical studies have shown that surface charging, internal charging, and electrostatic-discharge-related anomalies depend on ambient plasma and

energetic particle conditions, but also on spacecraft materials, illumination state, geometry, and the timescales over which charge accumulates and relaxes [5,14–16]. Evidence from geosynchronous communication spacecraft demonstrated that internal dielectric charging can produce recurring operational anomalies under enhanced energetic electron conditions [17,18]. Statistical studies of the GEO charging environment further showed that charging-relevant conditions vary systematically with plasma properties and orbital context [4]. Case studies of major GEO anomalies, including Galaxy 15 and Telstar 401, also indicated that unusual combinations of solar wind, geomagnetic, and particle conditions can create environments favorable for charging-related effects [6,7,19]. Together, this literature suggests that recurring anomalies may arise from persistent environmental forcing, geometry-dependent exposure, and finite-timescale spacecraft response.

The geostationary electron environment provides an important bridge between external forcing, local-time structure, and anomaly occurrence. Observations during the January 1997 magnetic storm showed that relativistic electron enhancements at geosynchronous orbit can vary strongly with local time and can be linked to magnetospheric compression and storm-time transport processes [19]. Earlier prediction-oriented studies demonstrated that geosynchronous electron environments can be modeled from in situ magnetic field information, highlighting the operational value of representing time-dependent environmental states [20]. More recent measurements from GOES and LANL platforms further characterized the variability of low- and medium-energy electrons at GEO and their dependence on local time and magnetospheric driving [21,22]. Predictive studies have also shown that geostationary electron fluxes can be modeled with useful skill when magnetic-local-time dependence is represented explicitly [23]. These findings are important for PHM-oriented anomaly analysis because they show that a local-time signal in anomaly occurrence should not automatically be interpreted as a purely geometric effect; it may also encode systematic variation in the surrounding plasma and charging environment.

Recent anomaly surveys and reviews have argued that spacecraft anomaly analysis should move beyond highly aggregated catalogs and should account for the orbit regime, anomaly type, subsystem context, and physical mechanism [8,9]. This direction is especially relevant for operational meteorological satellites, where space environment disturbances can affect service continuity and data availability [24]. From a PHM perspective, subsystem-level anomaly modes are closer to the actual units of monitoring, diagnosis, and operational response than mission-wide anomaly counts. General aerospace PHM research has similarly emphasized interpretable health-state assessment, fault diagnosis, and prognosis for complex systems operating under variable conditions [25]. Therefore, a recurring spacecraft anomaly mode should be analyzed not only as a retrospective space weather association, but also as a possible observable signature of an anomaly-prone system state.

The GOES-R series provides a useful setting for such subsystem-level PHM analysis because it combines long-duration geostationary operation, environmental monitoring, and well-documented instrument systems [10]. The Extreme Ultraviolet and X-ray Irradiance Sensors (EXIS) provide key solar X-ray and ultraviolet irradiance measurements for operational space weather monitoring [11]. Most directly related to the present study, Kress et al. [12] analyzed GOES-R operational anomalies together with in situ 30 eV–3 MeV electron measurements and showed that EXIS Space Wire anomalies are associated with enhanced electron fluxes and display recognizable local-time structure. This result establishes EXIS Space Wire anomalies as a focused recurring subsystem-level anomaly class, but it also raises a broader PHM question: whether their timing structure reflects spatial modulation, temporal clustering, external forcing, or the interaction of all three.

Despite these advances, several gaps remain in the PHM-oriented characterization of recurring GEO anomaly modes. First, many studies identify associations between anoma-

lies and space weather conditions, but fewer explicitly distinguish orbital–illumination modulation from external environmental forcing. This distinction is necessary because local-time and seasonal patterns may reflect both spacecraft geometry and the spatial organization of the magnetospheric environment. Second, temporal clustering is often observed or implied, but recent anomaly history is rarely treated as a separate state descriptor that must be distinguished from direct physical drivers. For PHM, this distinction is important because recent anomalies may indicate that the spacecraft–environment system remains in a susceptible or anomaly-active state, even when the underlying cause is not fully observed. Third, many analyses emphasize retrospective association within the available record, whereas PHM applications require evidence that the identified timing structure can generalize to later mission periods. Without time-held-out validation, it is difficult to determine whether a model has captured a temporally generalizable anomaly-state representation or only fitted period-specific correlations.

This study addresses these gaps by treating GOES-16/17 EXIS Space Wire anomalies as a validation case for PHM-oriented anomaly-state characterization. Rather than modeling anomaly timing as a single undifferentiated response, we decompose it into three complementary sources of information: orbital–illumination geometry, same-satellite event history structure, and external space weather forcing. We then evaluate these components within a satellite-stratified case–control framework and test their generalization using chronological train–test validation. In this way, the study extends previous GEO anomaly, charging environment, and GOES-R anomaly research toward a compact, interpretable, and temporally generalizable representation of subsystem-level anomaly-active states.

3. Proposed Method

The proposed framework is designed to convert a recurring anomaly catalog into an event-level anomaly-state characterization problem. It is built around a common physical time reference. Each anomaly or control sample is represented by its actual UTC timestamp, and all auxiliary information is aligned to that timestamp. Orbital and illumination variables are computed from satellite ephemeris and solar geometry because the local charging environment depends on where the spacecraft is located relative to the Sun, Earth, and magnetospheric structure. External space weather variables are matched using the most recent available prior OMNI observations and rolling histories because geomagnetic and solar wind forcing can persist over several hours rather than acting only instantaneously. Event history variables are constructed only from previous anomalies on the same satellite, so that they describe the recent anomaly-active state of the spacecraft–environment system without using future information. This time-aligned representation allows anomaly and non-anomaly samples to be compared under the same geometric, temporal, and environmental definitions.

3.1. Data Sources and Data Processing

The anomaly record used in this study is the public NOAA GOES-16/17 EXIS Space Wire anomaly timestamp data set. The file contains two satellite-specific sheets listing the UTC dates and times of recorded EXIS Space Wire anomalies for GOES-16 and GOES-17. During preprocessing, all timestamps were converted to UTC, records were sorted by satellite and time, and duplicate entries were checked before feature construction.

The retained anomaly record contains 1773 EXIS Space Wire anomaly events, including 919 events from GOES-16 and 854 events from GOES-17. The GOES-16 anomaly record extends from 10 December 2016 01:09:15.677 UTC to 14 March 2022 23:49:33.008 UTC, while the GOES-17 record extends from 28 March 2018 13:23:38.620 UTC to 14 March 2022 15:08:28.766 UTC. These anomaly timestamps were used as case observations in the

event-level case–control analysis. Table 1 summarizes the retained anomaly record by satellite. The overall analysis workflow is summarized in Figure 2.




Figure 2. Overview of the proposed analysis workflow. GOES-16/17 EXIS Space Wire anomaly timestamps are used as event anchors. Orbital and illumination descriptors are derived from satellite ephemeris and solar geometry at the corresponding UTC times, event history descriptors are constructed from previous same-satellite anomalies, and OMNI space weather variables are matched using prior-time environmental information. The resulting satellite-stratified case–control data set is evaluated through chronological train–test validation and clean selected-feature modeling.

Table 1. Summary of the retained GOES-16/17 EXIS Space Wire anomaly record.

Satellite	Start Time (UTC)	End Time (UTC)	Number of Events
GOES-16	10 December 2016 01:09:15.677	14 March 2022 23:49:33.008	919
GOES-17	28 March 2018 13:23:38.620	14 March 2022 15:08:28.766	854
Total	10 December 2016 01:09:15.677	14 March 2022 23:49:33.008	1773

Three groups of auxiliary data were then integrated with the anomaly record. First, GOES-R ephemeris and derived solar geometry quantities were used to describe orbital position, local solar time, illumination state, eclipse-season geometry, and seasonal solar geometry. These variables represent the spatial and illumination context in which an anomaly or control sample occurred. Second, same-satellite event history variables were computed from previously recorded anomalies to describe recent anomaly-state persistence. Third, hourly OMNI space weather data were matched to each sample to represent external geomagnetic, solar wind, interplanetary magnetic field (IMF), solar activity, and particle environment conditions. Exploratory lunar-geometric variables were also computed and included in the full candidate feature pool, but they were not used as the primary basis for physical interpretation unless retained by the selected-feature models.

3.2. Case–Control Design

The objective of the case–control design is to determine whether recorded anomaly times differ systematically from sampled non-anomaly times in orbital–illumination, event history, and space weather feature spaces. Each recorded anomaly timestamp was treated as a case, and non-anomaly control times were sampled from the same satellite-specific mission interval. Sampling controls separately for GOES-16 and GOES-17 preserves satellite-specific operating periods and avoids mixing different mission timelines.

For each anomaly case, five control timestamps were sampled, resulting in a nominal 1:5 case–control ratio. Candidate control times falling within ± 1 h of any recorded anomaly on the same satellite were rejected. This exclusion avoids labeling times immediately adjacent to known anomalies as controls while remaining narrow enough to reduce artificial

inflation of longer-timescale event history effects. The exclusion rule was applied only to control sampling; all anomaly cases were retained.

The resulting analysis table contains 10,638 samples, including 1773 anomaly cases and 8865 controls. The 1:5 ratio was preserved for each satellite: 919 GOES-16 cases were paired with 4595 controls, and 854 GOES-17 cases were paired with 4270 controls. The same feature construction procedure was applied to cases and controls, allowing sampled anomaly and non-anomaly times to be compared under consistent geometric, temporal, and environmental definitions.

3.3. Feature Construction

The candidate feature set was organized into four conceptual groups: orbital–illumination descriptors, event history descriptors, external space weather descriptors, and exploratory lunar-geometric descriptors. The full feature table contained 205 candidate predictors before feature screening, including 204 numerical predictors and one categorical satellite identifier.

3.3.1. Orbital and Illumination Features

Orbital and illumination features were designed to represent the spatial context of each sample. The satellite identifier was included to capture baseline differences between GOES-16 and GOES-17 after other predictors were considered. Local solar time was represented using harmonic terms to preserve its circular structure and avoid discontinuity between 24 h and 0 h. For sample i , local solar time LT_i was transformed as

$$\theta_i = 2\pi \frac{LT_i}{24}, \quad (1)$$

and the corresponding harmonic features were

$$\sin(\theta_i), \quad \cos(\theta_i), \quad \sin(2\theta_i), \quad \cos(2\theta_i). \quad (2)$$

Illumination was represented using the solar elevation angle at the spacecraft sub-satellite longitude, together with low-illumination threshold indicators. These variables were included because charging-related conditions and thermal/illumination states can change near low-sun and shadow-related geometries. Additional geometry and seasonal variables were generated, including satellite longitude, satellite radius, Earth-centered position variables, Sun–Earth–satellite angles, Sun–satellite–Earth angles, subsolar longitude, solar declination, days to nearest equinox, eclipse-season indicators, day-of-year harmonics, approximate beta-angle variables, Earth-shadow flags, and selected interaction terms.

3.3.2. Event History Features

Event history features were computed separately for each satellite using only anomaly information strictly prior to the sample time. Thus, the recent anomaly state of GOES-16 did not affect GOES-17 features, and vice versa. This design treats event history as a same-satellite state descriptor rather than as a mission-wide count.

The elapsed time since the most recent previous anomaly was used as a primary event history descriptor. For sample time t_i and previous anomaly times $\{t_j^{(a)}\}$ from the same satellite, the elapsed time was defined as

$$H_i = t_i - \max\{t_j^{(a)} : t_j^{(a)} < t_i\}, \quad (3)$$

and modeled as

$$x_i^{(\text{gap})} = \log(1 + H_i). \quad (4)$$

Trailing anomaly counts were also computed over 1 h, 3 h, 6 h, and 24 h windows, together with binary indicators for whether at least one anomaly occurred in selected recent windows. To describe smoother recent activity, an exponential decay intensity score was computed as

$$I_i^{(12h)} = \sum_{j: 0 < t_i - t_j^{(a)} \leq 72h} \exp\left[-\frac{t_i - t_j^{(a)}}{12h}\right], \quad (5)$$

with the corresponding transformed covariate

$$x_i^{(\text{int})} = \log\left(1 + I_i^{(12h)}\right). \quad (6)$$

Here, t_i denotes the timestamp of sample i , $t_j^{(a)}$ denotes the timestamp of the j -th previous anomaly on the same satellite, $I_i^{(12h)}$ denotes the exponentially decayed near-history indicator with a 12 h time constant, and the transformed covariate $x_i^{(\text{int})}$ provides a log-scale summary of recent anomaly intensity over the preceding 72 h. The constants 12 h and 72 h therefore define the decay scales of the short- and long-history descriptors, respectively.

This construction gives greater weight to recent anomalies while allowing weaker influence from earlier anomalies within the preceding 72 h. In the present study, event history variables are interpreted as descriptors of recent anomaly-active state rather than as direct evidence that one anomaly physically causes the next.

3.3.3. External Space Weather Features

External space weather features were constructed from hourly OMNI data. For each case or control sample, the most recent available OMNI observation prior to the sample time was used, and rolling-window summaries were computed over preceding 3 h, 6 h, 12 h, 24 h, and 72 h intervals. This design reflects the fact that geomagnetic and solar wind disturbances may influence the GEO environment over finite durations rather than only at the exact event time.

The geomagnetic activity variables included Kp, ap, Dst, and AE. Rolling maxima and means were computed for activity-enhancement variables such as Kp, ap, and AE, while rolling minima were also computed for quantities such as Dst and IMF B_z to represent storm-level or southward-field intervals. Binary storm indicators were included to mark intervals in which Kp reached storm-level conditions or Dst dropped below a storm threshold.

Solar wind and IMF predictors included solar wind speed, density, dynamic pressure, magnetic field magnitude, plasma beta, and IMF components in GSE and GSM coordinates. F10.7 solar radio flux was included as a solar activity background proxy, and energetic proton flux channels were included as particle environment indicators. These variables provide external forcing information that complements the local orbital–illumination and same-satellite event history descriptors.

3.3.4. Exploratory Lunar-Geometric Features

Exploratory lunar-geometric variables were computed to test whether non-solar geometry contributed to the timing structure. These variables included lunar distance, lunar phase angle, lunar illumination fraction, lunar declination, sublunar longitude and latitude, and Moon–Earth–satellite angular quantities. They were included in the full candidate feature pool but were not used as a central physical explanation unless selected by the clean feature-screening procedure. Table 2 summarizes the candidate feature groups and representative variables used in the extended analysis.

Table 2. Summary of candidate feature groups used in the extended analysis.

Feature Group	Representative Variables
Orbital geometry	Satellite identifier, satellite longitude, satellite radius, ECEF/GSE position variables, Sun–Earth–satellite angles, Sun–satellite–Earth angles
Local time and illumination	Local solar time harmonics, subsatellite-point solar elevation, low-sun indicators, very-low-sun indicators, Earth-shadow flag
Seasonal geometry	Solar declination, subsolar longitude, days to nearest equinox, eclipse-season flag, day-of-year harmonics, approximate beta-angle variables
Event history	Elapsed time since previous anomaly, trailing anomaly counts over 1 h, 3 h, 6 h, and 24 h, recent-event flags, burst ratio, 12 h exponential decay event intensity
Geomagnetic activity	Kp, ap, Dst, AE, storm-level flags, rolling maxima, minima, and means over preceding 3–72 h windows
Solar wind and IMF	Solar wind speed, density, dynamic pressure, plasma beta, IMF magnitude and components, especially GSM B_z rolling means and minima
Solar and particle environment	F10.7 solar radio flux and energetic proton flux channels with rolling-window summaries
Lunar geometry	Lunar distance, phase angle, illumination fraction, lunar declination, sublunar position, and Moon–Earth–satellite angular variables

3.4. Modeling and Evaluation

The sampled response variable was defined as a binary indicator. For observation i , $Y_i = 1$ denotes an anomaly case and $Y_i = 0$ denotes a control. Logistic regression was used to model the sampled case probability:

$$\Pr(Y_i = 1 \mid \mathbf{x}_i) = \pi_i, \quad (7)$$

with

$$\text{logit}(\pi_i) = \log\left(\frac{\pi_i}{1 - \pi_i}\right) = \beta_0 + \mathbf{x}_i^\top \boldsymbol{\beta}. \quad (8)$$

Because the analysis is based on a case–control sample, the fitted probabilities represent case probabilities within the sampled table rather than calibrated anomaly rates over continuous mission time. The main purpose of the models is therefore to compare the relative information provided by different feature groups and to evaluate whether this information generalizes to a later mission period.

To test temporal generalization, samples before 1 January 2021 were assigned to the training set, and samples from 1 January 2021 onward were assigned to the test set. The training set contains 7734 samples, including 1289 anomaly cases and 6445 controls. The test set contains 2904 samples, including 484 anomaly cases and 2420 controls. This chronological split allows model performance to be evaluated on a later portion of the anomaly record that was not used for fitting, preprocessing parameter estimation, or feature screening. Table 3 summarizes the chronological train–test split used for model evaluation.

All preprocessing steps that could learn from the data were fitted only on the training set and then applied to the test set. These steps included median imputation of continuous predictors, standardization of numerical features, categorical encoding of the satellite identifier, and L1-regularized feature screening. This procedure prevents information from the held-out period from influencing preprocessing or feature selection.

Table 3. Chronological train–test split used for time-held-out model evaluation.

Set	Satellite	Total Samples	Cases	Controls	Time Range
Train	GOES-16	4350	725	3625	10 December 2016 to 28 December 2020
Train	GOES-17	3384	564	2820	28 March 2018 to 30 December 2020
Train	All	7734	1289	6445	10 December 2016 to 30 December 2020
Test	GOES-16	1164	194	970	2 January 2021 to 14 March 2022
Test	GOES-17	1740	290	1450	12 January 2021 to 14 March 2022
Test	All	2904	484	2420	2 January 2021 to 14 March 2022

A sequence of logistic-regression models was used to evaluate the incremental contribution of each information source:

- **Orbital–Illumination Model (OIM):** This includes satellite identity, local-time harmonics, solar elevation variables, and selected orbital–illumination descriptors.
- **Event History Model (EHM):** This includes same-satellite event history descriptors, especially elapsed time since the previous anomaly and the 12 h exponential decay recent-activity score.
- **Orbital–Illumination–History Model (OIHM):** This combines the compact orbital–illumination and event history feature groups.
- **Space Environment-Integrated Model (SEIM):** This adds external space weather descriptors, including geomagnetic indices, solar wind parameters, IMF components, dynamic pressure, F10.7, and energetic particle proxies.
- **Clean Selected-feature models (CS models):** These use train-only L1-regularized logistic regression to screen the full candidate feature pool, followed by compact Top-15, Top-20, and Top-25 logistic models after removing direct near-neighbor event history variables.

This model sequence supports a layered interpretation of anomaly timing. The OIM tests whether anomaly occurrence contains persistent orbital and illumination modulation. The EHM tests whether the anomaly record contains temporal clustering consistent with anomaly-active states. The OIHM tests whether spatial and temporal structures provide complementary information. The SEIM tests whether external space weather forcing adds information beyond geometry and event history. The CS models test whether the full feature pool can be reduced to a compact and physically interpretable PHM-oriented feature representation.

For clean selected-feature modeling, L1-regularized logistic regression was fitted only on the training set. The retained features were ranked by the absolute values of their standardized L1 coefficients. Because the control sampling procedure excluded controls within ± 1 h of recorded anomalies, very-short-range event history descriptors may partly reflect the sampling design rather than generalizable anomaly-state information. The final clean Top- k models therefore removed direct near-neighbor history indicators, especially those describing anomaly occurrence or anomaly counts within the preceding 1 h. Closely related short-range binary indicators over the 3 h and 6 h windows were also excluded. This filtering step encouraged the compact models to rely on smoother anomaly-state descriptors, orbital–illumination information, and external space weather variables rather than immediate adjacency to previous anomalies.

Model discrimination was evaluated using the area under the receiver operating characteristic curve (AUC). The primary metric is the time-held-out test AUC computed on samples from 1 January 2021 onward. Training AUC values are reported as supporting diagnostics to show the in-sample fit, while the main interpretation is based on held-out

performance, incremental feature group comparison, and the physical composition of the selected predictors.

3.5. Baseline and Null Model

A constant intercept-only baseline (null model) was included to confirm that the reported discrimination is not a trivial consequence of class prevalence in the case–control sample. This model predicts p_{train} , the training-set case prevalence, as a constant score for all test samples, yielding an expected AUC of approximately 0.5.

3.6. Bootstrap Confidence Intervals and Paired Model Comparisons

To quantify uncertainty in the held-out metrics, stratified bootstrap 95% confidence intervals were computed for each model's test AUC, average precision (AP), and Brier score. Positive and negative samples were resampled separately within each bootstrap replicate to preserve class proportions, following the recommendation that bootstrap CIs for AUC should account for the case–control sampling structure. The default number of bootstrap replicates was 2000. Paired bootstrap AUC differences were used to compare models on the same held-out test set. For each pair (model A, model B), the difference $\text{AUC}(A) - \text{AUC}(B)$ was computed on the same bootstrap resamples, and a two-sided bootstrap p -value was obtained from the empirical distribution of differences. Significance was interpreted cautiously: when the 95% CI for the AUC difference included zero, the two models were described as providing comparable discrimination.

3.7. Exclusion Window Sensitivity Analysis

To assess whether the main conclusions depend on the adopted ± 1 h control exclusion rule, the entire case–control construction, external-feature integration, and model evaluation were repeated for exclusion windows of ± 0.5 h, ± 1 h, ± 3 h, and ± 6 h. The same random seed, controls-per-event ratio, chronological split date, and model specifications were used across all windows, so that differences reflect only the exclusion rule.

3.8. No-History Ablation and Drop-Column Group Ablation

To address the concern that event history descriptors might dominate the discrimination because they partly capture the temporal clustering induced by the exclusion window design, we added two complementary ablation strategies. First, no-history models (geometry-only, space weather-only, and geometry + space weather without any history variables) were evaluated. These quantify whether geometry and space weather variables retain discriminative value when all event history descriptors are removed. Second, drop-column group ablation was performed using the full model (geometry + history + space weather) as the reference: geometry, history, and space weather feature groups were dropped one at a time, and the resulting drop in held-out AUC was computed. Paired bootstrap comparisons between the full model and each ablated model were used to assess whether the AUC difference was statistically distinguishable from zero.

3.9. Group Permutation Importance

As an additional predictive diagnostic, group-level permutation importance was evaluated on the held-out test set. After fitting the full model on the training set, the columns belonging to each feature group (geometry, event history, space weather) were randomly permuted in the test set while keeping other columns unchanged. The resulting drop in test AUC, averaged over 100 independent permutations, was recorded. This measure reflects the loss in discriminative performance when the within-group feature structure is disrupted, and is interpreted as a predictive diagnostic rather than a causal attribution.

3.10. Cross-Satellite Generalization

To test whether the learned anomaly-state representation generalizes beyond within-mission temporal holdout, cross-satellite experiments were performed: training on GOES-16 and testing on GOES-17, and vice versa. In these experiments, the satellite-identifier dummy variable was excluded because it becomes constant in single-satellite training splits and offers no transferable information. Cross-satellite evaluation provides a stricter assessment than chronological holdout because it tests generalization across different spacecraft, operating periods, and possibly different local-environment conditions.

3.11. Calibration Diagnostic and Interpretation Limits

The case-control sampling design deliberately over-represents anomaly cases relative to their true operational frequency (1:5 ratio in the present study). Consequently, the fitted logistic regression scores do not represent calibrated operational anomaly probabilities. To provide an internal diagnostic, we computed reliability diagrams (binned observed vs. predicted scores) and fitted a logistic recalibration model $\text{logit}(Y) = \alpha + \beta \cdot \text{logit}(\hat{p})$ on the held-out test predictions. Under perfect calibration, $\alpha \approx 0$ and $\beta \approx 1$ would be expected for the case-control sample; however, the intercept α is not expected to be zero because the case-control prevalence does not match the operational prevalence. The slope β provides a measure of relative ordering quality. These diagnostics are reported as case-control internal assessments only, and the fitted scores are referred to throughout the manuscript as “anomaly-state scores” rather than as calibrated probabilities.

4. Results

4.1. Descriptive Characteristics

The retained anomaly record contains 1773 GOES-16/17 EXIS Space Wire anomaly events, including 919 events from GOES-16 and 854 events from GOES-17. These events form the case observations for the satellite-stratified case-control data set, which contains 10,638 total samples after control sampling. Figure 3 summarizes the main descriptive characteristics of the anomaly record and the derived analysis table.

The monthly-count distribution shows that anomaly occurrence varies over the mission interval rather than appearing as a uniform sequence of isolated events. The local-solar-time distribution provides an initial indication of orbital-phase dependence, while the solar elevation distribution reflects the illumination conditions represented in the case-control table. The elapsed-time distribution further shows that many samples occur at relatively short intervals after previous same-satellite anomalies, motivating the explicit inclusion of event history descriptors. Together, these descriptive views support the need to analyze EXIS Space Wire anomalies as a coupled spatial-temporal process rather than as independent event timestamps alone.

4.2. Time-Held-Out Model Comparison

Figure 4 summarizes the chronological train-test model comparison, and Table 4 reports the corresponding train and test AUC values. Models were trained on samples before 1 January 2021 and evaluated on samples from 1 January 2021 onward. The test AUC therefore measures whether the timing structure learned from the earlier mission period remains informative in a later held-out period.




Figure 3. Descriptive characteristics of the GOES-16/17 EXIS Space Wire anomaly record and derived case-control data. The panels show temporal occurrence, local-solar-time structure, illumination conditions, and the transformed elapsed-time feature used to represent recent anomaly-state history. (a) Monthly anomaly counts. (b) Local solar time at anomaly occurrence. (c) Solar elevation distribution in the case-control sample. (d) Distribution of $\log(1 + \text{hours since previous anomaly})$.




Figure 4. Summary of the time-held-out model comparison and feature-screening results. Panel (a) compares held-out test AUC values across sequential and clean selected-feature model configurations. Panel (b) shows feature retention from the 205-feature candidate pool to the train-only L1-retained set and the clean Top- k models. Panel (c) compares train and test AUC values for the clean Top-15, Top-20, and Top-25 selected-feature models. Panel (d) summarizes the feature group composition of the clean Top-15 model.

Table 4. Chronological train–test comparison of model discrimination. Models were trained on samples before 1 January 2021 and evaluated on samples from 1 January 2021 onward. AUC point estimates are the original chronological test set results. Stratified bootstrap 95% confidence intervals (2000 resamples) were added in the revised manuscript. The constant null baseline predicting the training-set case prevalence yields AUC = 0.500, confirming that the reported discrimination is not a trivial consequence of class prevalence.

Model	Features	Train AUC	Test AUC	Test AUC 95% CI
Orbital–Illumination Model (OIM)	6	0.5928	0.5978	[0.566, 0.616]
event history Model (EHM)	2	0.7316	0.6886	[0.658, 0.719]
Orbital–Illumination–History Model (OIHM)	8	0.7410	0.7133	[0.685, 0.740]
Space Environment–Integrated Model (SEIM)	30	0.8029	0.7540	[0.734, 0.783]
Clean Top-15, no direct 1 h near-neighbor variables	15	0.8060	0.7588	[0.725, 0.785]
Clean Top-20, no direct 1 h near-neighbor variables	20	0.8095	0.7544	[0.720, 0.780]
Clean Top-25, no direct 1 h near-neighbor variables	25	0.8120	0.7510	[0.715, 0.775]

The model sequence reveals a clear incremental pattern. The OIM reaches a test AUC of 0.5978, indicating that orbital geometry and illumination contain reproducible but limited timing information. This result shows that local time, illumination, and related orbital descriptors contribute to anomaly discrimination, but they are not sufficient as standalone predictors.

The EHM reaches a higher test AUC of 0.6886, showing that recent same-satellite anomaly history contains stronger information than geometry alone. This result supports the presence of temporal clustering in the EXIS Space Wire anomaly record. The OIHM further increases the test AUC to 0.7133, indicating that orbital–illumination information and recent anomaly-state history provide complementary contributions.

The SEIM achieves a test AUC of 0.7540 after adding external space weather variables. This improvement over the OIHM shows that geomagnetic activity, solar wind conditions, IMF structure, dynamic pressure, solar activity background, and particle environment proxies contribute additional held-out information beyond orbital–illumination and event history descriptors. Stratified bootstrap 95% confidence intervals (Table 4) provide uncertainty estimates for these model comparisons and help avoid over-interpreting small differences between closely performing models. The model sequence therefore supports a layered interpretation: EXIS Space Wire anomaly timing contains spatial modulation, temporal clustering, and externally forced environmental structure.

Adding lunar-geometric variables to the SEIM did not materially change held-out AUC (SEIM + Lunar: 0.756 [0.730, 0.780]; paired bootstrap $p \approx 1.0$). Among eight lunar candidates, four survived train-only L1 screening but none entered the clean Top-15 model.

Among the selected-feature models, the clean Top-15 model achieved the highest held-out AUC of 0.7588 after excluding direct 1 h near-neighbor history variables. This value was slightly higher than the full SEIM test AUC of 0.7540, but paired bootstrap comparisons indicated that the small difference was not statistically significant. The clean Top-20 and Top-25 models achieved test AUC values of 0.7544 and 0.7510, respectively. The slight decline with additional selected predictors suggests that the most temporally generalizable timing signal is captured by a compact subset of physically coherent features rather than by a larger selected-feature set. The clean Top-15 model is therefore emphasized as a compact and interpretable selected-feature representation, not as a statistically superior model.

Exclusion window sensitivity. The exclusion window sensitivity analysis (Table 5) showed that the SEIM test AUC increases gradually with wider exclusion windows: 0.756 [0.730, 0.781] at ± 0.5 h, 0.759 [0.733, 0.784] at ± 1 h, 0.779 [0.754, 0.803] at ± 3 h, and 0.799 [0.775, 0.822] at ± 6 h. The relative ranking of the key models remained broadly stable across windows. The history-only model was more sensitive to the exclusion window, confirming that event history descriptors partly capture the temporal clustering induced by the sampling design. We therefore interpret these variables as discriminative descriptors

under the adopted case–control framework rather than as independent physical drivers. The full model remains consistently discriminative across all tested windows (Figure 5b).




Figure 5. Reviewer-requested validation and robustness analyses. (a) Time-held-out model comparison with stratified bootstrap 95% confidence intervals. (b) Exclusion window sensitivity analysis from ± 0.5 h to ± 6 h; the SEIM remains discriminative across all windows while the absolute AUC level varies with the sampling window. (c) Group-level ablation and permutation importance diagnostics quantifying the contributions of orbital–illumination geometry, event history descriptors, and space weather variables. (d) Cross-satellite generalization tests between GOES-16 and GOES-17, showing partial transfer with moderated interpretation. Asterisks denote paired-bootstrap significance levels: ** $p < 0.01$ and *** $p < 0.001$.

Table 5. Exclusion window sensitivity analysis. Test AUC values with stratified bootstrap 95% confidence intervals for four representative models. The SEIM remains discriminative across all tested windows, while the history-only model varies more with the exclusion rule.

Model	± 0.5 h	± 1 h	± 3 h	± 6 h
OIM (geometry-only)	0.596 [0.568, 0.624]	0.598 [0.569, 0.624]	0.598 [0.570, 0.624]	0.601 [0.572, 0.629]
EHM (history-only)	0.686 [0.656, 0.716]	0.689 [0.658, 0.719]	0.709 [0.679, 0.738]	0.732 [0.702, 0.761]
Geom + SW (no history)	0.745 [0.719, 0.770]	0.748 [0.722, 0.772]	0.762 [0.737, 0.786]	0.775 [0.750, 0.798]
SEIM (full model)	0.756 [0.730, 0.781]	0.759 [0.733, 0.784]	0.779 [0.754, 0.803]	0.799 [0.775, 0.822]

Bold indicates the full SEIM reference model.

Cross-satellite generalization. Cross-satellite tests provided a stricter assessment than chronological holdout. Training on G16 and testing on G17 yielded a SEIM test AUC of 0.806 [0.789, 0.824], while training on G17 and testing on G16 yielded 0.770 [0.753, 0.787] (Figure 5d). Both directions showed discrimination well above chance, supporting partial cross-satellite transfer. The asymmetry between the two directions indicates that the representation contains both shared environmental susceptibility information and satellite- or period-specific structure.

Calibration diagnostics. Calibration analysis of the SEIM on the held-out test set yielded a logistic recalibration slope of $\beta = 0.972$ ($p = 0.582$ for $H_0 : \beta = 1$), indicating the adequate relative ordering of the discriminative scores. The intercept was $\alpha = -1.756$ ($p < 0.001$ for $H_0 : \alpha = 0$), confirming the expected prevalence-related bias under case–control sampling. These results support the interpretation that the model outputs should

be treated as discriminative anomaly-state scores rather than as calibrated operational anomaly probabilities.

4.3. Clean Selected-Feature Model

The full candidate feature pool contained 205 predictors, including orbital-geometry variables, local-time harmonics, illumination and eclipse-season descriptors, event history variables, geomagnetic indices, solar wind and IMF variables, solar activity and energetic particle proxies, and exploratory lunar-geometric variables. Train-only L1-regularized logistic regression retained 71 nonzero features. These retained features were ranked by the absolute value of their standardized L1 coefficients.

Because the control sampling design excluded controls within ± 1 h of recorded anomalies, direct near-neighbor history variables were removed from the final clean selected-feature models. This step prevents the compact models from depending primarily on immediate adjacency to previous anomalies and encourages them to retain broader anomaly-state, orbital-illumination, and space environment information.

Table 6 lists the predictors retained in the clean Top-15 model. The selected set is dominated by external space weather variables, especially solar wind density, IMF B_z , AE, and Kp. Orbital-illumination information remains present through local-time harmonics, solar declination, and Earth-shadow state. Event history is represented by the elapsed time since the previous same-satellite anomaly, rather than by direct 1 h near-neighbor indicators.

Table 6. Clean Top-15 selected features after train-only L1 screening and removal of direct 1 h near-neighbor history variables. Coefficients are standardized L1 logistic coefficients from the training-set screening step.

Rank	Feature	Feature Group	Standardized Coefficient
1	Solar wind density mean (24 h)	Space weather	−0.1875
2	IMF B_z GSM mean (72 h)	Space weather	−0.1663
3	Local-time cosine	Orbital-illumination	−0.1574
4	Log time since previous anomaly	Event history	−0.1516
5	AE max (6 h)	Space weather	0.1501
6	AE current	Space weather	0.1487
7	Solar declination	Orbital-illumination	0.1311
8	AE max (72 h)	Space weather	0.1167
9	Kp mean (24 h)	Space weather	0.1123
10	Local-time sine	Orbital-illumination	0.1110
11	Kp max (72 h)	Space weather	0.1098
12	AE mean (6 h)	Space weather	0.1025
13	Kp max (24 h)	Space weather	0.1000
14	Earth-shadow flag	Orbital-illumination	0.0980
15	IMF B_z GSM mean (6 h)	Space weather	0.0924

Figure 6 visualizes the standardized coefficient directions of the clean Top-15 model. The coefficient signs describe conditional associations within the standardized training design. They should be read together with the selected feature groups rather than as isolated physical effects, because rolling-window space weather variables are often correlated.

The retained space weather variables indicate that external forcing is a major component of the compact timing signature. AE and Kp variables contribute positive coefficients, linking higher sampled anomaly-case odds with enhanced auroral/substorm activity and broader geomagnetic disturbance. IMF B_z and solar wind density variables also remain among the highest-ranked predictors, indicating that upstream forcing and magnetospheric coupling conditions contribute information over multi-hour to multi-day windows.

The retained orbital-illumination variables indicate that the compact model does not reduce anomaly timing to space weather alone. Local-time sine and cosine terms preserve orbital-phase dependence, while solar declination and Earth-shadow state represent seasonal and illumination-related modulation. The event history component is captured by

the log time since the previous anomaly, whose negative coefficient indicates higher sampled anomaly-case odds at shorter elapsed times after previous same-satellite anomalies. Because direct 1 h near-neighbor variables were excluded, this result supports a broader anomaly-active state interpretation rather than simple immediate adjacency.




Figure 6. Standardized L1 logistic coefficients for the clean Top-15 selected-feature model. Coefficients are from train-only L1 screening after median imputation and standardization. Positive standardized coefficients increase the log-odds of being classified as an anomaly case under the case-control design, whereas negative coefficients decrease it, conditional on the other standardized predictors. Because several retained predictors are correlated, coefficient signs and magnitudes should be interpreted as conditional associations rather than direct causal effects. Direct 1 h near-neighbor history variables were excluded from the clean Top-*k* models.

Overall, the clean Top-15 model provides a compact representation of the main timing structure identified in the full analysis. Its predictors span three physically meaningful categories: external space weather forcing, orbital-illumination modulation, and conservative event history information. This composition is consistent with the model comparison results and supports the interpretation that the GOES-16/17 EXIS Space Wire record contains a coupled spatial-temporal-environmental anomaly-state signature suitable for PHM-oriented monitoring.

5. Discussion

5.1. Interpretation of the Coupled Anomaly-State Signature

The results support the central premise of this study: recurring spacecraft anomalies should not be interpreted only as isolated event outcomes or as responses to a single external driver. Instead, the GOES-16/17 EXIS Space Wire anomaly record contains a coupled anomaly-state signature that combines orbital-illumination modulation, temporal clustering, and external space weather forcing. This finding is important for PHM because it suggests that the observed timing structure carries information about when the spacecraft-environment system becomes more susceptible to anomalies.

Orbital-illumination variables provide reproducible but limited discrimination. The Orbital-Illumination Model indicates that local time, seasonal geometry, and shadow-related conditions contribute to anomaly timing, but their standalone predictive power is

modest. This result is consistent with previous studies showing that geostationary anomaly occurrence and charging-relevant electron environments vary with local time and orbital context [3,4,21–23]. However, the limited discrimination also shows that geometry should not be treated as a complete explanation of the anomaly process. In the PHM interpretation, orbital–illumination context is better viewed as a background modulation factor that changes the susceptibility of the spacecraft–environment system rather than as the sole cause of anomaly occurrence.

The stronger performance of the event history Model indicates that EXIS Space Wire anomalies exhibit substantial temporal clustering. This clustering is not merely a descriptive feature of the event sequence; it is central to anomaly-state characterization. A short elapsed time since previous same-satellite anomalies may indicate that the spacecraft remains in an anomaly-active interval. Such intervals may arise from persistent geomagnetic disturbance, sustained energetic particle enhancement, charging accumulation and relaxation, subsystem recovery behavior, or interactions between environmental and engineering states. Therefore, recent anomaly history should be interpreted as an observable state descriptor rather than as direct evidence that one anomaly physically triggers another. This distinction is important for PHM because event history can be operationally useful even when the underlying physical mechanism is only partially observed.

Adding external space weather variables further improves held-out discrimination, showing that the anomaly-state signature is not fully captured by temporal clustering alone. The Space Environment-Integrated Model improves upon the combined orbital–illumination and history model, indicating that geomagnetic activity, solar wind conditions, IMF structure, dynamic pressure, solar activity background, and particle environment proxies provide complementary information. The clean selected-feature model reinforces this interpretation. AE and Kp represent substorm and geomagnetic activity, IMF B_z reflects magnetospheric coupling conditions, solar wind density describes upstream forcing, and local-time, solar-declination, and Earth-shadow variables represent geometric and illumination context. The retained elapsed-time feature connects these external and geometric factors to recent anomaly-state persistence.

Taken together, these results show that the EXIS Space Wire anomaly sequence is best understood as a temporally clustered anomaly-state process modulated by both orbital–illumination context and external space weather forcing. The key point is not that one factor dominates all others, but that each factor contributes a different layer of information. Geometry describes where and under what illumination conditions the spacecraft is operating, space weather describes the external forcing environment, and event history summarizes whether the system has recently exhibited anomaly-active behavior. Their combination provides a more complete and physically interpretable representation than any single component alone.

5.2. Implications for PHM-Oriented Anomaly-State Monitoring

From a PHM perspective, the main contribution of this study is the decomposition of anomaly timing into interpretable and temporally generalizable components. A monitoring strategy based only on static orbital geometry would miss much of the observed temporal and environmental structure. A strategy based only on recent anomaly occurrence would identify active episodes but would provide limited physical context. A strategy based only on space weather indices would ignore the fact that the same external forcing may have different implications depending on orbital–illumination context and recent subsystem behavior. The results therefore support a PHM representation that combines orbital–illumination modulation, recent anomaly-state history, and external space weather forcing.

The clean Top-15 model is especially relevant for this interpretation. After excluding direct 1 h near-neighbor history variables, the compact model retained the best held-out performance while using a physically coherent set of predictors. This suggests that the generalizable timing signal does not depend primarily on immediate adjacency to previous anomalies. Instead, it can be represented by a compact combination of geomagnetic activity, solar wind and IMF variables, local-time and illumination descriptors, and elapsed time since the previous anomaly. Such a feature set is valuable for PHM because it is not only predictive in a sampled case–control sense, but also interpretable as a representation of the spacecraft–environment state.

The results also refine the interpretation of previous GOES-R anomaly studies. Prior work showed that GOES-R operational anomalies, including EXIS Space Wire anomalies, are associated with enhanced electron fluxes and display local-time structure [12]. The present study builds on that finding by showing that the local-time signal should not be interpreted as purely geometric, and that temporal clustering should not be treated merely as a sampling artifact. Both are embedded in a broader anomaly-state structure that includes external forcing and recent system behavior. This provides a more complete subsystem-level basis for operational anomaly monitoring.

More broadly, the proposed framework illustrates how recurring anomaly records can be used for PHM even when detailed engineering telemetry is unavailable. By aligning anomaly timestamps, sampled control times, orbital–illumination descriptors, event history variables, and space weather conditions within a common event-level design, the framework converts an anomaly catalog into a structured anomaly-state analysis problem. This approach can support screening of temporally generalizable timing signatures, identification of physically meaningful monitoring variables, and prioritization of future telemetry or environmental data needed for mechanism-level diagnosis.

5.3. Interpretation Boundaries

The results of this study should be understood within clearly stated interpretation boundaries. We distinguish three levels of inference that are often conflated in anomaly analysis.

First, **predictive relevance** refers to whether a variable or feature group improves held-out discrimination. The bootstrap CIs, ablation analyses, and permutation importance reported here support the predictive relevance of event history, space weather, and, to a lesser extent, orbital–illumination descriptors under the case–control design.

Second, **statistical association** refers to whether a predictor is conditionally associated with the case–control label in a multivariable model. The standardized coefficients and odds ratios reported for the final models reflect conditional associations within the sampled design matrix and should not be interpreted as independent physical effects. Because several space weather predictors (e.g., AE, Kp, and their rolling-window summaries) are substantially correlated (Spearman $|r|$ up to 0.92), individual coefficient signs and magnitudes may reflect shared variability rather than isolated variable-level effects. Figure 7 provides calibration and correlation diagnostics supporting the interpretation limits discussed below.

Third, **physical causality** would require additional physical modeling, mechanism-level evidence from in situ measurements or engineering telemetry, or prospective validation under controlled conditions. The present study does not establish causal relationships between specific space weather variables and EXIS Space Wire anomaly occurrence. The retained variables should therefore be interpreted as predictive descriptors of the anomaly-state score rather than as direct causal drivers of the anomalies.

In addition, the case–control design with a ± 1 h exclusion window increases the apparent discriminative value of event history descriptors, as confirmed by the exclusion

window sensitivity analysis. event history variables should be interpreted as descriptors of temporal clustering under the adopted sampling design, not as independent physical causes. The no-history ablation models further show that space weather variables retain discriminative value after removing all history descriptors, although their standalone performance is lower than that of the full model.




Figure 7. Calibration and feature correlation diagnostics. **(a)** Case–control calibration diagnostic (reliability curve and recalibration statistics) showing that fitted scores should be interpreted as discriminative anomaly-state indicators rather than calibrated operational probabilities. **(b)** Spearman correlation structure among retained predictors, motivating the interpretation of standardized coefficients as conditional multivariable associations rather than independent causal effects.

5.4. Limitations and Future Work

Several limitations should be considered when interpreting the results. First, the case–control framework is designed to compare recorded anomaly times with sampled non-anomaly times. The fitted scores therefore represent sampled case discrimination rather than calibrated anomaly occurrence rates over continuous mission time. Calibration diagnostics confirmed that the intercept is not zero ($\alpha < 0$), consistent with the expected case–control prevalence bias, while the slope was reasonably close to 1 ($\beta \approx 0.972$), indicating adequate relative ordering. The chronological train–test split and cross-satellite tests strengthen the evidence for temporal and inter-satellite generalization, but a fully operational PHM system would require validation on continuous mission timelines and calibration of event probabilities, event rates, or hazard scores through prospective time-grid validation or prevalence correction.

Second, the external space weather variables used in this study are primarily OMNI-based indicators and rolling-window summaries. These variables describe upstream solar wind and geomagnetic forcing, but they do not fully represent the local plasma and electron environment at each spacecraft. Future work should incorporate detailed in situ electron measurements across charging-relevant energy ranges, local magnetic-time information, and spacecraft-specific environmental measurements. This would help connect the statistical anomaly-state signature more directly with physical charging mechanisms.

Third, the present analysis does not include detailed spacecraft engineering telemetry or subsystem operational state. Such data would be necessary to distinguish whether temporal clustering reflects environmental persistence, subsystem recovery behavior, repeated interface-level susceptibility, operational mode dependence, or combined environmental–engineering effects. Incorporating telemetry, command history, thermal state, operational mode, and recovery actions would allow a more complete PHM interpretation of anomaly-active states.

Fourth, the cross-satellite generalization tests showed partial transfer of the learned representation. While discrimination was stronger than chance, it was weaker than within-mission chronological generalization in some directions. This indicates that the representation contains both shared environmental susceptibility information and satellite- or period-specific structure. The original term “transferable anomaly-state signature” has therefore been moderated throughout the revised manuscript in favor of more cautious terminology such as “temporally generalizable anomaly-state representation” and “candidate anomaly-state signature.”

Fifth, among the eight lunar-geometric candidate variables, four survived train-only L1 screening, but none entered the clean selected-feature models. This result suggests that lunar geometry did not provide stable incremental information after accounting for the retained orbital-illumination, event history, and space weather variables. Adding lunar variables to the full model did not materially improve held-out AUC ($\Delta\text{AUC} < 0.005$, paired bootstrap $p \approx 1.0$). We therefore treat lunar descriptors as evaluated but not robustly retained predictors in the present data set.

Finally, the feature-screening process necessarily deals with correlated predictors, especially among rolling-window summaries of space weather variables. Individual coefficient signs should therefore be interpreted as conditional model associations rather than isolated physical effects. The most robust conclusions rest on the consistency of the model sequence, the held-out performance gains, the ablation and permutation results, and the repeated selection of physically meaningful feature groups. Future work should evaluate point-process models, survival or hazard-based formulations, continuous-time PHM monitoring schemes, and multi-mission validation to test whether the same spatial-temporal-environmental decomposition generalizes to other recurring spacecraft anomaly modes.

6. Conclusions

Advantages and main findings. This study developed a case-control discriminative framework for characterizing recurring GOES-16/17 EXIS Space Wire anomaly states using orbital-illumination geometry, event history descriptors, and space weather variables. Chronological holdout evaluation showed that orbital-illumination geometry alone provides weak but reproducible discrimination (test AUC ≈ 0.60), event history descriptors capture stronger temporal clustering (test AUC ≈ 0.69), and space weather variables add complementary information beyond both geometry and history, with the full SEIM reaching a held-out test AUC of 0.7540. A compact train-only L1-selected Top-15 model achieved comparable held-out performance after excluding direct near-neighbor history variables, retaining a physically interpretable combination of AE, Kp, IMF B_z , solar wind density, local-time harmonics, solar declination, Earth-shadow state, and elapsed time since the previous anomaly. Stratified bootstrap 95% confidence intervals and paired bootstrap comparisons provided quantitative uncertainty estimates, and the constant null baseline confirmed that the reported discrimination is not a trivial consequence of class prevalence.

Limitations. The case-control design does not support the direct interpretation of fitted scores as calibrated operational anomaly probabilities; calibration diagnostics confirmed the expected prevalence-related intercept bias. The exclusion window sensitivity analysis showed that event history descriptors are partly influenced by the ± 1 h sampling rule, and their discriminative contribution should be interpreted as descriptive of temporal clustering under the adopted design rather than as independent physical driving. Cross-satellite generalization tests showed partial transfer, indicating that the learned representation contains both shared environmental susceptibility information and satellite- or period-specific structure. Several space weather predictors were strongly correlated, so individual coefficient signs and magnitudes should be interpreted as conditional multivariable

associations rather than as direct causal effects. Lunar-geometric variables were evaluated, but none entered the clean selected-feature models. The analysis does not include in situ electron measurements or engineering telemetry, which would be needed to connect the statistical anomaly-state signature more directly to physical charging mechanisms.

Future perspectives. Future work should pursue prospective validation on a continuous operational time grid with appropriate prevalence correction to enable calibrated anomaly probability or hazard estimation. Incorporating in situ electron flux measurements across charging-relevant energy ranges, local magnetic-time information, and spacecraft-specific environmental data would strengthen the connection between the anomaly-state representation and physical charging mechanisms. Integrating engineering telemetry, command history, thermal state, and operational mode would support a more complete PHM interpretation. Extending the framework to additional satellites, instruments, and anomaly classes would test the generality of the decomposition across different spacecraft and operational contexts. Online anomaly-state monitoring using the compact selected-feature representation, possibly within a point-process or survival-analysis formulation, represents a practical PHM application pathway.

Author Contributions: Conceptualization, T.Y. and D.G.; methodology, T.Y. and Z.L.; software, T.Y.; validation, T.Y., Z.L. and Y.L.; formal analysis, T.Y. and Z.L.; investigation, T.Y.; resources, D.G., F.D. and X.L.; data curation, T.Y. and Y.L.; writing—original draft preparation, T.Y. and Z.L.; writing—review and editing, T.Y., Z.L., D.G., F.D., Y.L. and X.L.; visualization, T.Y.; supervision, D.G. and F.D.; project administration, D.G. and X.L.; funding acquisition, D.G., F.D. and X.L. All authors have read and agreed to the published version of the manuscript.

Funding: This research was partially funded by the National Natural Science Foundation of China, grant number 42371476.

Institutional Review Board Statement: Not applicable.

Informed Consent Statement: Not applicable.

Data Availability Statement: The data used in this study are publicly available. The GOES-16/17 EXIS Space Wire anomaly records are available from NOAA public data resources. GOES-R ephemeris and derived orbital information are available from public GOES-R data products. The OMNI space weather data are available from the NASA OMNIWeb database. The processed case-control table and analysis scripts can be made available by the authors upon reasonable request.

Acknowledgments: The authors acknowledge the providers of the public GOES-R and OMNI space weather data used in this study.

Conflicts of Interest: The authors declare no conflicts of interest. The funder had no role in the design of the study; in the collection, analyses, or interpretation of data; in the writing of the manuscript; or in the decision to publish the results.

Abbreviations

The following abbreviations are used in this manuscript:

AE	Auroral electrojet index
AP	Average precision
AUC	Area under the receiver operating characteristic curve
CI	Confidence interval
CS	Clean selected-feature model
EHM	Event history model
EXIS	Extreme ultraviolet and X-ray irradiance sensors
GEO	Geostationary earth orbit
GOES	Geostationary operational environmental satellite

IMF	Interplanetary magnetic field
Kp	Planetary geomagnetic index
L1	Least absolute shrinkage and selection operator/L1 regularization
OIM	Orbital–illumination model
OIHM	Orbital–illumination–history Model
OMNI	NASA OMNI space weather data set
PHM	Prognostics and health management
ROC	Receiver operating characteristic
SEIM	Space environment-integrated model
SpW	SpaceWire

List of Symbols

The following symbols are used in this manuscript:

t	Sample timestamp (UTC)
Y	Case–control label ($Y = 1$: anomaly case; $Y = 0$: control)
$\mathbf{x}$	Feature vector
$\hat{p}$	Fitted anomaly-state score (case–control discriminative score)
π	Sampled case probability under the case–control design
β_0	Logistic regression intercept
$\boldsymbol{\beta}$	Logistic regression coefficient vector
AUC	Area under the receiver operating characteristic curve
ΔAUC	AUC difference between two paired models
h	Exclusion window half-width in hours
τ	Exponential decay time constant for event history score (12 h)
H	Elapsed time since previous same-satellite anomaly
$I^{(12h)}$	12 h exponential decay event intensity score
θ	Local solar time angle (radians)
α	Calibration intercept (logistic recalibration)
β	Calibration slope (logistic recalibration)

References

- Wu, J.-G.; Eliasson, L.; Lundstedt, H.; Hilgers, A.; Andersson, L.; Norberg, O. Space environment effects on geostationary spacecraft: Analysis and prediction. *Adv. Space Res.* **2000**, *26*, 31–36. [https://doi.org/10.1016/S0273-1177\(99\)01023-6](https://doi.org/10.1016/S0273-1177(99)01023-6).
- Pilipenko, V.; Yagova, N.; Romanova, N.; Allen, J. Statistical relationships between satellite anomalies at geostationary orbit and high-energy particles. *Adv. Space Res.* **2006**, *37*, 1192–1205. <https://doi.org/10.1016/j.asr.2005.03.152>.
- Choi, H.-S.; Lee, J.; Cho, K.-S.; Kwak, Y.-S.; Cho, I.-H.; Park, Y.-D.; Kim, Y.-H.; Baker, D.N.; Reeves, G.D.; Lee, D.-K. Analysis of GEO spacecraft anomalies: Space weather relationships. *Space Weather* **2011**, *9*, S06001. <https://doi.org/10.1029/2010SW000597>.
- Thomsen, M.F.; Henderson, M.G.; Jordanova, V.K. Statistical properties of the surface-charging environment at geosynchronous orbit. *Space Weather* **2013**, *11*, 237–244. <https://doi.org/10.1002/swe.20049>.
- Mazur, J.E.; Fennell, J.F.; Roeder, J.L.; O'Brien, P.T.; Guild, T.B.; Likar, J.J. The timescale of surface-charging events. *IEEE Trans. Plasma Sci.* **2012**, *40*, 237–245. <https://doi.org/10.1109/TPS.2011.2174656>.
- Loto'aniu, T.M.; Singer, H.J.; Rodriguez, J.V.; Green, J.; Denig, W.; Biesecker, D.; Angelopoulos, V. Space weather conditions during the Galaxy 15 spacecraft anomaly. *Space Weather* **2015**, *13*, 484–502. <https://doi.org/10.1002/2015SW001239>.
- Saiz, E.; Cid, C.; Guerrero, A. Environmental conditions during the reported charging anomalies of the two geosynchronous satellites: Telstar 401 and Galaxy 15. *Space Weather* **2018**, *16*, 1784–1801. <https://doi.org/10.1029/2018SW001974>.
- Buitrago-Leiva, J.N.; El Khayati Ramouz, M.; Camps, A.; Ruiz-de-Azua, J.A. Statistical analysis of LEO and GEO satellite anomalies and space radiation. *Aerospace* **2024**, *11*, 924. <https://doi.org/10.3390/aerospace11110924>.
- Minow, J.I.; Jordanova, V.K.; Pitchford, D.; Ganushkina, N.Y.; Zheng, Y.; Delzanno, G.L.; Jun, I.; Kim, W. ISWAT spacecraft surface charging review. *Adv. Space Res.* **2024**, *in press*. <https://doi.org/10.1016/j.asr.2024.08.058>.
- Goodman, S.J.; Schmit, T.J.; Daniels, J.M.; Redmon, R.J. GOES-R series summary and look ahead. In *The GOES-R Series: A New Generation of Geostationary Environmental Satellites*; Goodman, S.J., Schmit, T.J., Daniels, J.M., Redmon, R.J., Eds.; Elsevier: Amsterdam, The Netherlands, 2020; pp. 273–277. <https://doi.org/10.1016/B978-0-12-814327-8.00023-8>.

11. Machol, J.L.; Eparvier, F.G.; Viereck, R.A.; Woodraska, D.L.; Snow, M.; Thiemann, E.; Woods, T.N.; McClintock, W.E.; Mueller, S.; Eden, T.D., Jr.; et al. GOES-R series solar X-ray and ultraviolet irradiance. In *The GOES-R Series: A New Generation of Geostationary Environmental Satellites*; Goodman, S.J., Schmit, T.J., Daniels, J.M., Redmon, R.J., Eds.; Elsevier: Amsterdam, The Netherlands, 2020; pp. 233–242. <https://doi.org/10.1016/B978-0-12-814327-8.00019-6>.
12. Kress, B.T.; Rodriguez, J.V.; Buzulukova, N.Y.; Redmon, R.J.; Machol, J.L.; Fiorello, J.L.; Roza, M.A.; Meloy, R.M. Relationship between GOES-R series spacecraft operational anomalies and in situ 30 eV–3 MeV electron measurements. *IEEE Trans. Plasma Sci.* **2024**, *52*, 1610–1618. <https://doi.org/10.1109/TPS.2024.3390658>.
13. O'Brien, T.P. SEAES-GEO: A spacecraft environmental anomalies expert system for geosynchronous orbit. *Space Weather* **2009**, *7*, S09003. <https://doi.org/10.1029/2009SW000473>.
14. Garrett, H.B.; Whittlesey, A.C. Spacecraft charging, an update. *IEEE Trans. Plasma Sci.* **2000**, *28*, 2017–2028. <https://doi.org/10.1109/27.902229>.
15. Matéo-Vélez, J.-C.; Sicard, A.; Payan, D.; Ganushkina, N.; Meredith, N.P.; Sillanpää, I. Spacecraft surface charging induced by severe environments at geosynchronous orbit. *Space Weather* **2018**, *16*, 89–106. <https://doi.org/10.1002/2017SW001689>.
16. Ganushkina, N.Y.; Swiger, B.; Dubyagin, S.; Matéo-Vélez, J.-C.; Liemohn, M.W.; Sicard, A.; Payan, D. Worst-case severe environments for surface charging observed at LANL satellites as dependent on solar wind and geomagnetic conditions. *Space Weather* **2021**, *19*, e2021SW002732. <https://doi.org/10.1029/2021SW002732>.
17. Wrenn, G.L. Conclusive evidence for internal dielectric charging anomalies on geosynchronous communications spacecraft. *J. Spacecr. Rocket.* **1995**, *32*, 514–520. <https://doi.org/10.2514/3.26645>.
18. Wrenn, G.L.; Rodgers, D.J.; Ryden, K.A. A solar cycle of spacecraft anomalies due to internal charging. *Ann. Geophys.* **2002**, *20*, 953–956. <https://doi.org/10.5194/angeo-20-953-2002>.
19. Reeves, G.D.; Friedel, R.H.W.; Belian, R.D.; Meier, M.M.; Henderson, M.G.; Onsager, T.; Singer, H.J.; Baker, D.N.; Li, X.; Blake, J.B. The relativistic electron response at geosynchronous orbit during the January 1997 magnetic storm. *J. Geophys. Res. Space Phys.* **1998**, *103*, 17559–17570. <https://doi.org/10.1029/97JA03236>.
20. Tsutai, A.; Mitsui, C.; Nagai, T. Prediction of a geosynchronous electron environment with in situ magnetic field measurements. *Earth Planets Space* **1999**, *51*, 219–223. <https://doi.org/10.1186/BF03352225>.
21. Sillanpää, I.; Ganushkina, N.Y.; Dubyagin, S.; Rodriguez, J.V. Electron fluxes at geostationary orbit from GOES MAGED data. *Space Weather* **2017**, *15*, 1602–1614. <https://doi.org/10.1002/2017SW001698>.
22. van de Kamp, M.; Ganushkina, N.; Simms, L.; Liemohn, M. Drivers for geostationary 2–200 keV electron fluxes as observed at GOES satellites. *Space Weather* **2024**, *22*, e2024SW003984. <https://doi.org/10.1029/2024SW003984>.
23. Simms, L.E.; Ganushkina, N.Y.; van de Kamp, M.; Liemohn, M.W. Predicting geostationary (GOES) 4.1–30 keV electron flux over all MLT using LEEMYR regression models. *Space Weather* **2024**, *22*, e2024SW003962. <https://doi.org/10.1029/2024SW003962>.
24. Sakaguchi, K.; Nagatsuma, T. Impact of space environment on geostationary meteorological satellite data outage. *Space Weather* **2022**, *20*, e2021SW002965. <https://doi.org/10.1029/2021SW002965>.
25. Nguyen, V.D.; Kefalas, M.; Yang, K.; Apostolidis, A.; Olhofer, M.; Limmer, S.; Bäck, T. A review: Prognostics and health management in automotive and aerospace. *Int. J. Progn. Health Manag.* **2019**, *10*, 35. <https://doi.org/10.36001/ijphm.2019.v10i2.2730>.

Disclaimer/Publisher's Note: The statements, opinions and data contained in all publications are solely those of the individual author(s) and contributor(s) and not of MDPI and/or the editor(s). MDPI and/or the editor(s) disclaim responsibility for any injury to people or property resulting from any ideas, methods, instructions or products referred to in the content.