Available online at www.sciencedirect.com

ScienceDirect

Advances in Space Research xx (2026) xxx-xxx

ADVANCES IN SPACE RESEARCH (a COSPAR publication)

www.elsevier.com/locate/asr

# Complexity and scaling descriptors as diagnostic predictors of heliophysical indices across solar-cycle timescales

D. Sierra-Porta[a,*], Maximiliano Canedo Verdugo[b], Daniel David Herrera Acevedo[a]

[a] *Universidad Tecnológica de Bolívar. Escuela de Transformación Digital., Parque Industrial y Tecnológico Carlos Vélez Pombo Km 1 Vía Turbaco., Cartagena de Indias, 130010, Colombia*

[b] *Universidad de Sonora. Departamento de Física., Blvd. Colosio y Calle de la Sabiduría, Colonia Centro. Edificio 3F., Hermosillo, 83000, México*

Received —; Received in final form —; Accepted —;
Available online —



## Abstract

Heliophysical variability emerges from a coupled, multiscale system in which changes in the solar atmosphere and heliospheric plasma translate into measurable signatures in widely used activity indices. Operational space-weather workflows often summarize this variability through amplitudes and a small set of bulk solar-wind covariates, yet important dynamical information may also reside in the evolving *morphology* of the signals. We examine whether shape descriptors computed from heliophysical time series provide information beyond classical amplitude summaries and standard bulk solar-wind covariates. Using daily OMNIWeb-era records spanning 1964–2025, we compute ten sliding-window descriptors under a past-only convention, designed to capture complementary aspects of temporal morphology such as irregularity, roughness, and long-range dependence. The descriptor set combines entropy measures, fractal-dimension estimators, the Hurst exponent, and Lempel–Ziv (LZ) complexity, yielding a compact representation of time-series structure that is not reducible to amplitude alone. The window length is treated as a methodological hyperparameter and selected through a target-specific sensitivity analysis that jointly favors competitive out-of-sample RMSE and stable permutation-importance rankings across neighboring windows.

Two complementary learners, gradient boosting and a multilayer perceptron, are used as diagnostic probes to quantify permutation-based feature relevance under chronological splitting and training-only preprocessing. Across three targets (F10.7, Sunspot Number, and Dst), shape descriptors consistently rank among the most informative predictors, often matching or exceeding the relevance of standard solar-wind inputs. The most robust signals arise from LZ complexity and a compact subset of entropy/fractal measures, whose windowed trajectories track solar-cycle phases with characteristic lead–lag behaviour. Correlation analyses on both levels and standardised first differences expose redundancy within descriptor families and reduce spurious associations driven by shared nonstationarity, motivating a family-level interpretation of relevance rather than causal attribution. Overall, the results indicate that heliophysical time-series morphology encodes dynamical information complementary to amplitude- and bulk-plasma descriptions, suggesting compact, instrument-light features for augmenting future space-weather modelling pipelines.



*Keywords:* Heliospheric medium ; Solar radio flux ; Sunspots ; Complexity ; signal-shape descriptors
*2000 MSC:* 28A80 ;
*2000 MSC:* 28D20 ; 37B40
*2000 MSC:* 94A17
*2000 MSC:* 85A35
*2000 MSC:* 70K55 ; 85A35 ; 85A20 ; 86A10 ; 37A60

*Corresponding author: Email address: dporta@utb.edu.co (D. Sierra-Porta)
*Email address:* dporta@utb.edu.co (D. Sierra-Porta)

https://dx.doi.org/10.1016/j.jasr.xxxx.xx.xxx

## 1. Introduction

Topological Data Analysis (TDA) (Carlsson & Vejdemo-Johansson, 2021; Chazal & Michel, 2021) grew out of applied algebraic topology and computational geometry and has popularized the idea that the *shape* of data can carry information complementary to standard amplitude-based summaries. In heliophysics and space-weather research, geometry-aware characterizations have become increasingly relevant for long, heterogeneous time series where dynamical regimes and multiscale variability coexist (Gopinath & Prince, 2017; Sierra-Porta, 2022; Sarlis et al., 2024; Sierra-Porta & Domínguez-Monterroza, 2022). Formal topological descriptors have also been explored in related solar applications—e.g., Betti-number based feature sets for flare prediction (Deshmukh et al., 2023)—highlighting the broader value of shape information, even though our descriptors are different in nature.

From a methodological standpoint, the TDA program has matured into scalable algorithms and data structures for extracting topological summaries from point-cloud representations, and multiple open-source libraries now provide robust implementations for persistent-homology pipelines (e.g., GUDHI and related ecosystems (Maria et al., 2014; Fasy et al., 2014; Pérez et al., 2021)). In this broader landscape, we focus on a complementary, time-series regime: long heliospheric records motivate fast, windowed complexity and scaling descriptors that can be computed daily and integrated seamlessly with standard statistical learning workflows.

Fractal and multifractal analyses have proven informative in space-physics observables beyond amplitude statistics; for instance, multifractal detrended fluctuation analysis has been used to link cosmic-ray intensities to geomagnetic cutoff rigidity (Sierra-Porta & Domínguez-Monterroza, 2022). At solar-cycle scales, multifractal characterizations have also revealed phase-dependent differences among sunspot number, magnetic-field strength, and irradiance-related observables, suggesting that the same activity level can correspond to distinct fluctuation geometries across datasets and phases (De Freitas & De Medeiros, 2009). These results are causing widespread concern that underlies this study: whether signal morphology carries dynamical information that is not reducible to amplitude-based descriptors.

The heliospheric medium, primarily consisting of solar-wind plasma and interplanetary magnetic fields, exhibits complex, non-equilibrium behavior in which intermittent structures, long-range correlations, and multiscale fluctuations are routinely observed. Such complexity is often better summarized by entropy- and scaling-based descriptors than by mean levels alone, particularly during Solar–Magnetospheric Interaction (SMI) events where non-extensive signatures and nonequilibrium transitions have been reported (Stumpo et al., 2020; Pavlos et al., 2019; Livadiotis, 2020). These studies suggest that heliospheric variability can involve self-organization processes, reduced effective dimensionality, and anomalous transport mechanisms that standard second-order statistics struggle to capture comprehensively.

A growing body of work supports the use of entropy, scaling, and complexity descriptors to discriminate heliospheric and geomagnetic regimes. For example, permutation-entropy and complexity measures, together with Hurst-exponent analyses, have been used to differentiate solar-wind structures at 1 AU, including slow and fast wind, magnetic clouds, ICME-driven sheaths, and stream interaction regions (SIRs), with magnetic clouds typically exhibiting the most ordered fluctuations (Kilpua et al., 2024). Detailed studies of ICME-driven sheath regions further indicate slightly higher complexity and lower entropy relative to preceding slow wind, consistent with their mixed composition of coherent structures and disordered fluctuations (Kilpua et al., 2022).

In geomagnetic indices, storm-time dynamics have likewise shown distinctive signatures, including changes in fractal dimension (and corresponding changes in Hurst-type persistence) and decreases in entropy measures during intense storms (Balasis et al., 2009; Wawrzaszek et al., 2022; Donner et al., 2018). These examples also highlight a practical issue for feature attribution: several families of descriptors are partially redundant, and theoretical links exist among persistence/roughness, fractal-dimension concepts, and spectral behavior, so importance can redistribute among correlated predictors (Gneiting & Schlather, 2004). Accordingly, rather than treating any single descriptor as a unique marker, we focus on robustness at the level of recurring feature families and stability under methodological perturbations (e.g., window length).

From a statistical perspective, the approach is also connected to earlier ideas in latent-parameter and pattern-theory modeling, where observables are augmented with summaries intended to capture hidden structure (Mumford, 2002). In nonlinear dynamics, related questions are often posed in terms of observability and embedding-based reconstruction from data (Aguirre & Letellier, 2011; Sauer et al., 1991). This viewpoint motivates our use of flexible supervised learners as function approximators, not primarily for operational forecasting but to obtain model-agnostic relevance estimates for shape descriptors relative to standard physical covariates.

This study builds on and generalises our cycle–scale survey, Global complexity signatures of solar cycles: A unified Entropy–Fractal survey of OMNI solar wind data (1964–2025) (Sierra-Porta et al., 2025). While that study identified two stable macro–axes of complexity (amplitude breadth and temporal irregularity) and documented their phase dependence using global, cycle–aggregated metrics, the present work operationalises the same framework at daily resolution. We move from descriptive, unsupervised characterisation (correlations, PCA, nonparametric tests) to supervised, model–agnostic attribution via permutation importance, and from global measurements to windowed signal-shape descriptors aligned with heliophysical indices. The window length is treated as a methodological choice and is selected empirically via sensitivity analysis (and may differ across target variables), rather than being fixed a priori. In this paper, however, we do not compute persistent-homology objects (e.g., persistence diagrams or Betti curves). Instead, we adopt a *TDA-inspired* viewpoint and work with computationally lightweight *signal-shape descriptors* derived from entropy, scaling, and complexity.

Our central goal is to test whether the morphology of heliophysical time series encodes dynamical information that is not reducible to amplitude-based characteristics. Using daily OMNIWeb records (IMF and plasma parameters) together with three target indices—F10.7, Sunspot Number, and Dst—we compute a suite of windowed signal-shape descriptors including multiple entropies, fractal-dimension estimates, the Hurst exponent, and Lempel–Ziv complexity. Lempel–Ziv complexity is of particular interest because of its links to algorithmic (Kolmogorov–Chaitin) complexity and its role as an estimator of dynamical/topological entropy in certain settings (Kaspar & Schuster, 1987; Brudno, 1978). We then use supervised learners (neural networks and gradient boosting) as flexible function approximators and assess feature relevance through permutation importance. Because time series can display spurious associations due to shared trends or oscillatory components, and because statistical dependence does not imply causality, we interpret these attributions as descriptive/diagnostic rather than causal (Yule, 1926). Together, these steps provide a practical, model-agnostic assessment of which aspects of signal shape most consistently explain variability in key heliophysical indices across solar-cycle timescales.

## 2. Data, materials and methods

### 2.1. Data used

We draw on NASA/GSFC's OMNIWeb average daily data set, which supplies homogenized heliospheric and geomagnetic measurements from 1964 (start of Solar Cycle 20) through 2025. The variables analysed are: (i) the magnitude of the interplanetary magnetic field (IMF); (ii) solar-wind bulk speed; (iii) proton temperature; (iv) proton number density; (v) the $\alpha$-to-proton ratio; and three diagnostic indices of solar and geomagnetic activity—Dst, Sunspot Number, and the 10.7 cm radio flux (F10.7). For completeness we add the Kp index to the list of variables even though it is correlated with the Dst index and additionally a secondary derived variable such as the solar-wind dynamic (flow) pressure $P_{\rm dyn} = \rho V^2 \simeq m_p\, n_p\, V^2$, (where $n_p$ is the proton number density, $V$ the bulk speed, and $m_p$ the proton mass). Although $P_{\rm dyn}$ is partly redundant—being mechanically determined by $n_p$ and $V$ and therefore correlated with them—it provides a compact proxy for the momentum flux of the solar wind that is commonly used in geospace analyses. This six-decade record spans Solar Cycles 20 through 25, capturing multiple minima and maxima and therefore a broad spectrum of heliophysical conditions.

Each series underwent standard quality control (removal of fill values and physically implausible points) followed by multivariate gap filling with a Bayesian linear Iterative Imputer. Overall, missing observations represent less than 28% of the total in every variable, with the $\alpha$/proton ratio exhibiting the highest fraction because of historically intermittent $\alpha$-particle measurements. Post-imputation checks confirmed that first- and second-order statistics (mean, variance) and inter-variable coherency were preserved, yielding a continuous, physically consistent data matrix ready for windowed signal-shape descriptor extraction (entropy-, scaling-, and complexity-based measures) and subsequent machine-learning experiments. To avoid information leakage in supervised analyses, all preprocessing steps that learn from data (including the imputer and any feature scaling) are fitted using the training split only and then applied to the held-out period.

### 2.2. Materials and methods

#### 2.2.1. Signal-shape descriptors

To probe the *shape* of each target index—Sunspot Number, the geomagnetic storm-time index (Dst), and the 10.7 cm solar radio flux (F10.7)—we extract ten well-established entropy-, scaling-, and complexity-based *signal-shape descriptors* from nonlinear-dynamics and information-theoretic literature. Although motivated by the broader shape-of-data viewpoint popularized by TDA, we therefore refer to them as *descriptors* rather than topological metrics.

Five descriptors are entropies that quantify uncertainty or disorder at different levels of representation. Shannon entropy (Shannon, 1948; Chen, 2021) summarizes the information content of the amplitude distribution; larger values indicate broader, less concentrated histograms. Sample entropy (Richman & Moorman, 2000; Richman et al., 2004) focuses on the conditional probability that short sequences of observations remain similar when extended, providing a noise-robust gauge of local irregularity. Permutation entropy (Bandt & Pompe, 2002) examines ordinal patterns (order relations of neighbouring points), offering a scale-aware view of complexity that is sensitive to dynamical changes even under monotone transformations. Spectral entropy treats the normalised power spectral density as a probability law, translating frequency-space disorder into a single summary statistic. Approximate entropy measures the likelihood that patterns that are close for $m$ points remain close for $m + 1$, thus capturing coarse regularity with low computational cost.

Three additional descriptors quantify curve roughness through fractal-dimension estimators. Higuchi, Katz, and Petrosian fractal dimensions (Higuchi, 1988; Petrosian, 1995) estimate, via different algorithms, the self-affine irregularity of a record; together they summarize how jagged or space-filling the trajectory appears across scales. Lempel–Ziv (LZ) complexity (Zozor et al., 2005) counts the rate at which new substrings appear in a symbolic rendering of the series, reflecting the novelty of evolving patterns. LZ is particularly relevant because it connects to algorithmic (Kolmogorov–Chaitin) complexity and has been related to dynamical/topological entropy in certain settings (Kaspar & Schuster, 1987; Brudno, 1978). Finally, the Hurst exponent $H$ (Carbone et al., 2004; Mielniczuk & Wojdyłło, 2007) summarizes long-range dependence: values near 0.5 are consistent with weakly correlated

increments, $H > 0.5$ indicates persistence, and $H < 0.5$ indicates antipersistence. In idealized self-affine models, $H$ is often linked to a graph (time-series) fractal dimension through relations of the form $D \approx 2 - H$, but such identities depend on modeling assumptions and should not be interpreted as universal equalities (Gneiting & Schlather, 2004). Related scaling diagnostics have been applied to heliophysical time series (e.g., multifractal DFA), reinforcing the relevance of roughness and memory beyond first-order moments (Sierra-Porta & Domínguez-Monterroza, 2022).

We emphasize that several descriptors are partially redundant by design: entropies and scaling measures can covary strongly, and theoretical connections between persistence and roughness imply correlated predictors. Rather than attributing unique causal meaning to any single descriptor, we later interpret relevance in terms of robust, recurring feature families and sensitivity to methodological choices such as window length. We also note that complementary structure-aware families—for example, visibility-graph and network-based indices—have been successfully applied to related heliophysical time series (Sierra-Porta, 2024); here we purposefully keep a minimalist panel to retain interpretability and parsimony.

Each descriptor is computed on sliding windows of the target series (Section 2.2.2), producing a time-aligned panel of shape features that complements the contemporaneous physical covariates. For a more detailed definition of these descriptors and their heliophysical interpretation, we refer the reader to Sierra-Porta et al. (2025) and the references therein.

### 2.2.2. *Recursive sliding-window scheme*

Rather than compressing the entire multi-decade record into a single value, we analyse each target series $x(t)$ at a daily cadence using a sliding window (Perea & Harer, 2015; Braverman, 2015) of length $w^*$ days and step $s = 1$ day. The window length is treated as a methodological hyperparameter and selected empirically via a sensitivity analysis; accordingly, we use target-specific values $w^*_{\mathrm{F10.7}} = 80$ days, $w^*_{\mathrm{SN}} = 60$ days, and $w^*_{\mathrm{Dst}} = 70$ days (Section 3.1). For a total length $L$ (days), the number of windows is $N = (L - w)/s + 1$, yielding a dense, day-by-day representation of local signal morphology over 1964–2025.

Let $x(t)$ denote a daily-sampled target index and let $t$ index days. For each day $t$ we compute ten signal-shape descriptors on the segment $x(t - w + 1), \ldots, x(t)$, i.e., a *past-only* window. This causal convention avoids using future observations when forming features and therefore prevents information leakage when the resulting design matrix is split chronologically. In practice, the first $w - 1$ days do not admit a full window and are discarded. The window length spans multiple solar-rotation periods (∼27 days) while retaining responsiveness to mesoscale geomagnetic variability.

For each target we thus obtain an $N \times 10$ matrix of windowed descriptors, $T^{(\mathrm{target})} = \{\mathbf{z}(t)\}_t$. In the supervised attribution experiments, we construct a target-specific design matrix by concatenating these shape descriptors with contemporaneous physical covariates from OMNIWeb,

$$\mathbf{X}^{(\mathrm{target})} = \left[\, T^{(\mathrm{target})} \,\middle|\, \mathbf{P} \,\right] \in \mathbb{R}^{N \times p},$$

where $\mathbf{P}$ contains the selected solar-wind variables (e.g., IMF magnitude $B$, plasma speed $V_p$, proton temperature $T_p$, proton density $N_p$, and the $\alpha$/p ratio). Crucially, we do not include the target amplitude $x(t)$ itself as a predictor, so that the models cannot trivially reconstruct the target from an identity feature. Each target (F10.7, Sunspot Number, and Dst) is analysed with its own $\mathbf{X}^{(\mathrm{target})}$ and response vector $\mathbf{y}^{(\mathrm{target})} = x(t)$.

### 2.2.3. *Feature normalisation and permutation-based variable-importance analysis*

To place every predictor on a comparable footing and avoid information leakage, we standardise all columns of the design matrix using training-only statistics under a chronological train/validation split. For each split we compute the mean $\mu_{\mathrm{train}}$ and standard deviation $\sigma_{\mathrm{train}}$ from the training period and apply them to both the training and held-out samples. This split-wise scaling is essential because the feature set combines physical magnitudes (e.g., Dst in nT) with dimensionless shape descriptors (entropies, fractal dimensions, complexities) whose numerical ranges differ substantially.

The normalised matrix feeds two multilayer perceptron (MLP) (Kruse et al., 2022; Popescu et al., 2009) configurations and a gradient-boosted tree regressor (XGBoost) (Chen et al., 2015; Ali et al., 2023). These models are not intended as operational forecasting systems; rather, they act as diagnostic probes that reveal how strongly each predictor contributes to explaining variability of the target series.

To quantify that contribution we apply the permutation-importance algorithm (Altmann et al., 2010) to every trained model. After the baseline error is recorded, each feature column is randomly permuted while all others are kept intact; the increase in prediction error (RMSE) is measured, and the procedure is repeated $R$ times (with independent random seeds) to obtain a stable ranking of relevance. Because permutation importance can redistribute attribution among strongly correlated predictors, we interpret the resulting rankings primarily in terms of robust, recurring feature families and assess their stability across window lengths and model classes (Gneiting & Schlather, 2004). This stability criterion is reported explicitly in the window-length sensitivity analysis (Section 3.1).

For XGBoost we additionally verify qualitative consistency against native gain and SHAP summaries, but we report permutation importance throughout for methodological uniformity across neural and tree-based learners. Although forecasting skill is not the objective here, stable permutation rankings provide an operationally useful screen for identifying which shape descriptors complement, or surpass, conventional solar-wind variables across solar-cycle timescales, thereby informing feature sets for future predictive pipelines.

## 3. Results

All results are reported for the three targets (F10.7, Sunspot Number, and Dst). The learning models are used as diagnostic probes to obtain permutation-based relevance rankings, not as operational forecasters. Windowed features are computed with a *past-only* convention, and the window length is selected empirically via sensitivity analysis, yielding target-specific values $w^*_{\mathrm{F10.7}} = 80$ days, $w^*_{\mathrm{SN}} = 60$ days, and $w^*_{\mathrm{Dst}} = 70$ days (Section 3.1).

### 3.1. Window-length sensitivity and stability of relevance rankings

We evaluate the robustness of model-based attribution to the sliding-window length by repeating the full pipeline across a grid of candidate windows. For each target and learner, we summarise (i) descriptive fit (RMSE) and (ii) stability of feature relevance using the mean Spearman correlation of permutation-importance vectors across windows and the mean top-$k$ overlap. We then select $w^*$ by prioritising stability (high Spearman/top-$k$) while requiring competitive RMSE, rather than simply choosing the minimum-RMSE window.

Figure 1 makes clear that the two diagnostics capture distinct aspects of the pipeline. RMSE varies smoothly with $w^*$ and often improves gradually as the window widens, reflecting the stronger noise suppression and the increased effective sample size of each descriptor estimate. In contrast, importance stability is markedly more structured: for each target, the Spearman stability curves exhibit well-defined plateaus and local maxima, indicating window ranges where the attribution pattern is reproducible across the tested grid. This separation between fit and stability is precisely why the window length must be treated as a methodological hyperparameter: selecting $w^*$ solely by RMSE can favour windows that fit the target slightly better while producing less reproducible rankings of which descriptors matter.

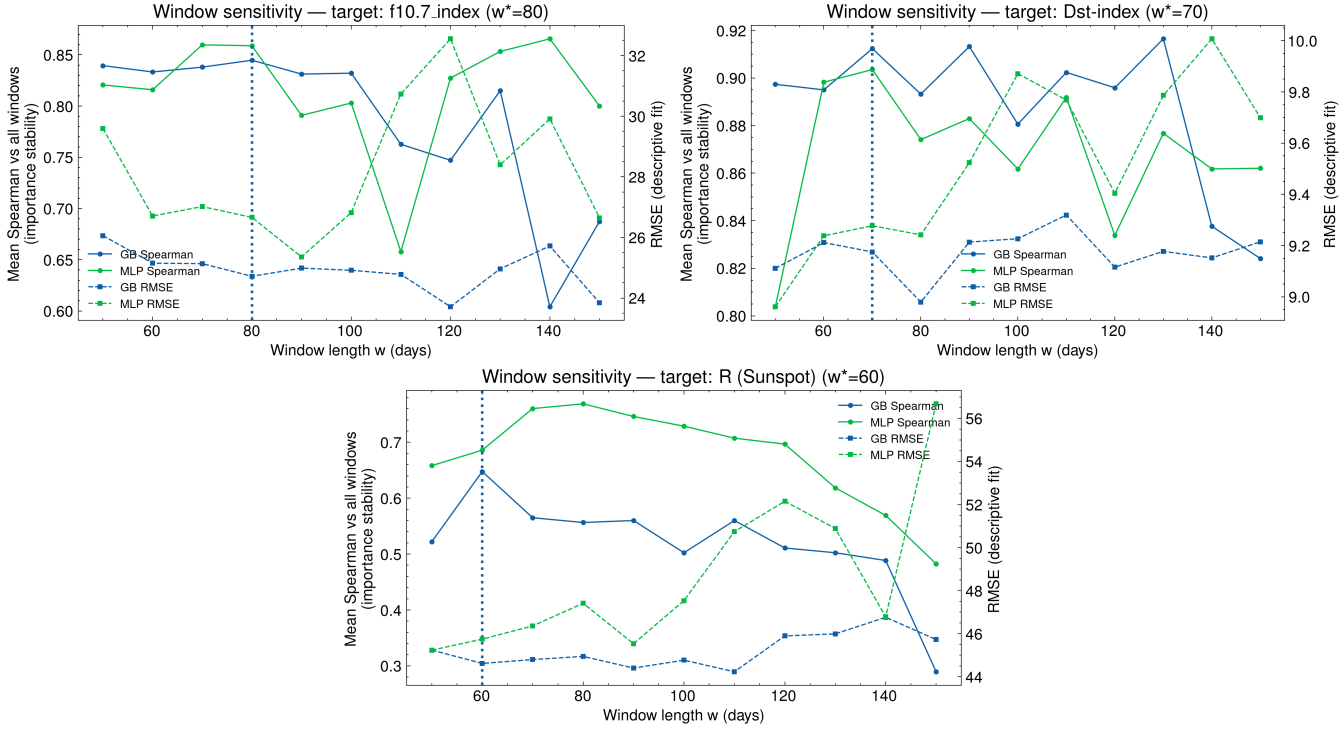


Fig. 1: Window-length sensitivity and stability of relevance rankings. Solid lines show the mean Spearman stability of permutation-importance vectors versus all tested windows; dashed lines show RMSE (descriptive fit). The vertical dashed line in each panel marks the selected target-specific window $w^*$.

The three panels in Fig. 1 further show that the optimal window depends on the target and on the dynamical timescales that dominate it. For Dst, stability is consistently high (Spearman $\gtrsim 0.88$ for both learners across most windows) with a pronounced maximum around $w^* = 70$ days, while RMSE differences across neighbouring windows remain modest. This behaviour suggests that the descriptor-derived representation of magnetospheric variability is robust to moderate changes in temporal context, and that $w^* \approx 70$ days offers a favourable trade-off between variance reduction and sensitivity to storm-time morphology. For F10.7, by contrast, stability remains strong up to $w^* \approx 100$ days but degrades at very long windows for at least one learner, consistent with the idea that overly wide windows can dilute the local morphological signatures that track the rising/declining phases of solar activity. Sunspot Number shows the lowest stability overall, which is expected given the discretised and intermittently bursty nature of daily sunspot counts; nevertheless, a clear stability peak occurs at $w^* = 60$ days, supporting a shorter context that preserves the episodic structure of active regions while still smoothing day-to-day counting noise.

Quantitatively, Table 1 confirms that the selected $w^*$ values deliver high stability (mean Spearman $\rho_S \approx 0.91$ for Dst and $\rho_S \approx 0.85$ for F10.7) with substantial top-$k$ agreement between windows (typically $\gtrsim 0.83$). While the RMSE-optimal window can differ from $w^*$ for some targets and learners, those gains are not accompanied by a commensurate increase in stability, and in some cases coincide with less consistent attributions. Because the central goal of this study is *explanatory attribution*—identifying which shape descriptors consistently carry information beyond classical parameters—we adopt $w^*$ as the window length that maximises reproducibility of relevance patterns under a competitive descriptive fit. All subsequent analyses therefore use the target-specific $w^*$ windows.

Table 1: Window-length sensitivity summary and selected target-specific windows $w^*$. Stability is quantified by the mean Spearman correlation of permutation-importance vectors against all tested windows and by the mean top-$k$ overlap; descriptive fit is reported as RMSE on the held-out split.

| **Target** | $w^*$ | $\rho_S$ | | | **Top-$k$** | | | **RMSE** | | |
|---|---|---|---|---|---|---|---|---|---|---|
| | | **(GB)** | **(MLP)** | **(mean)** | **(GB)** | **(MLP)** | **(mean)** | **(GB)** | **(MLP)** | **(mean)** |
| **Dst-index** | 70 | 0.912 | 0.904 | 0.908 | 0.89 | 0.89 | 0.890 | 9.175 | 9.278 | 9.227 |
| | 90 | 0.913 | 0.883 | 0.898 | 0.89 | 0.89 | 0.890 | 9.214 | 9.524 | 9.369 |
| | 110 | 0.902 | 0.892 | 0.897 | 0.85 | 0.90 | 0.875 | 9.320 | 9.769 | 9.544 |
| | 130 | 0.916 | 0.877 | 0.897 | 0.87 | 0.88 | 0.875 | 9.178 | 9.786 | 9.482 |
| | 60 | 0.895 | 0.898 | 0.897 | 0.88 | 0.87 | 0.875 | 9.212 | 9.239 | 9.225 |
| | 80 | 0.893 | 0.874 | 0.884 | 0.88 | 0.87 | 0.875 | 8.980 | 9.242 | 9.111 |
| | 100 | 0.881 | 0.862 | 0.871 | 0.85 | 0.86 | 0.855 | 9.227 | 9.871 | 9.549 |
| | 120 | 0.896 | 0.834 | 0.865 | 0.88 | 0.87 | 0.875 | 9.116 | 9.405 | 9.260 |
| | 50 | 0.897 | 0.804 | 0.851 | 0.83 | 0.87 | 0.850 | 9.111 | 8.961 | 9.036 |
| | 140 | 0.838 | 0.862 | 0.850 | 0.88 | 0.86 | 0.870 | 9.152 | 10.007 | 9.580 |
| **R (Sunspot)** | 60 | 0.648 | 0.686 | 0.667 | 0.79 | 0.87 | 0.830 | 44.620 | 45.748 | 45.184 |
| | 70 | 0.565 | 0.760 | 0.663 | 0.77 | 0.87 | 0.820 | 44.799 | 46.362 | 45.580 |
| | 80 | 0.556 | 0.769 | 0.663 | 0.75 | 0.87 | 0.810 | 44.946 | 47.412 | 46.179 |
| | 90 | 0.560 | 0.746 | 0.653 | 0.74 | 0.82 | 0.780 | 44.404 | 45.535 | 44.969 |
| | 110 | 0.560 | 0.707 | 0.634 | 0.74 | 0.83 | 0.785 | 44.232 | 50.741 | 47.487 |
| | 100 | 0.502 | 0.729 | 0.616 | 0.78 | 0.83 | 0.805 | 44.770 | 47.537 | 46.154 |
| | 120 | 0.511 | 0.697 | 0.604 | 0.76 | 0.82 | 0.790 | 45.897 | 52.150 | 49.024 |
| | 50 | 0.522 | 0.658 | 0.590 | 0.80 | 0.73 | 0.765 | 45.237 | 45.234 | 45.235 |
| | 130 | 0.502 | 0.619 | 0.560 | 0.73 | 0.72 | 0.725 | 45.992 | 50.889 | 48.440 |
| | 140 | 0.489 | 0.569 | 0.529 | 0.78 | 0.75 | 0.765 | 46.762 | 46.784 | 46.773 |
| **f10.7_index** | 80 | 0.845 | 0.859 | 0.852 | 0.85 | 0.90 | 0.875 | 24.719 | 26.662 | 25.690 |
| | 70 | 0.838 | 0.860 | 0.849 | 0.85 | 0.83 | 0.840 | 25.135 | 27.019 | 26.077 |
| | 130 | 0.815 | 0.853 | 0.834 | 0.85 | 0.90 | 0.875 | 24.965 | 28.397 | 26.681 |
| | 50 | 0.839 | 0.821 | 0.830 | 0.85 | 0.86 | 0.855 | 26.059 | 29.597 | 27.828 |
| | 60 | 0.833 | 0.816 | 0.825 | 0.81 | 0.90 | 0.855 | 25.150 | 26.707 | 25.929 |
| | 100 | 0.832 | 0.803 | 0.817 | 0.87 | 0.81 | 0.840 | 24.916 | 26.821 | 25.869 |
| | 90 | 0.831 | 0.791 | 0.811 | 0.87 | 0.77 | 0.820 | 24.988 | 25.356 | 25.172 |
| | 120 | 0.747 | 0.827 | 0.787 | 0.80 | 0.86 | 0.830 | 23.713 | 32.554 | 28.133 |
| | 150 | 0.687 | 0.800 | 0.744 | 0.75 | 0.86 | 0.805 | 23.843 | 26.648 | 25.246 |
| | 140 | 0.604 | 0.866 | 0.735 | 0.76 | 0.86 | 0.810 | 25.718 | 29.912 | 27.815 |

### *3.2. Correlation structure and descriptor redundancy*

To characterise redundancy among predictors, we examine rank correlations between the windowed shape descriptors and the physical covariates. Following classical cautions on "nonsense correlations" in time series (Yule, 1926), we report Spearman correlations both on raw levels and on standardised first differences (computed after windowing). The first-difference view reduces associations that can arise from shared slow trends, cycle-scale envelopes, or broad-band oscillatory components, and it therefore provides a more conservative picture of instantaneous co-variability.

Figure 2 shows the correlation matrices for the standardised first differences at the target-specific windows $w^*_{\mathrm{F10.7}} = 80$ days, $w^*_{\mathrm{SN}} = 60$ days, and $w^*_{\mathrm{Dst}} = 70$ days. (For completeness, the corresponding level-based matrices are provided in the Supplementary Material.) Across the three targets, the dominant structure is consistent: an *entropy–roughness* family groups the information-theoretic measures (Shannon, Sample, Permutation, Spectral, Approximate) together with the fractal-dimension estimators (Higuchi, Katz, Petrosian). Within this family, correlations are typically moderate-to-strong, indicating that these descriptors respond to overlapping aspects of short-scale irregularity and curve roughness. This is expected: several of these quantities are

mathematically or asymptotically linked through persistence/roughness relations and spectral scaling, so substantial collinearity is not a pathology but a natural consequence of measuring closely related properties with different estimators (Gneiting & Schlather, 2004). Here, we show first-difference correlations to attenuate common-trend effects and reduce spurious associations; level-based correlations are provided in the Supplementary Material for completeness.

A second, more weakly coupled group is formed by Lempel–Ziv complexity and the Hurst exponent. Their correlations with the entropy–roughness family are comparatively small and often of opposite sign, suggesting that algorithmic novelty and long-range dependence accentuate dynamical aspects that are not reducible to local disorder or multi-scale jaggedness. This observation aligns with the interpretation that LZ is closer to dynamical-complexity notions grounded in symbolic dynamics and entropy rates, whereas the entropy–roughness block is dominated by distributional/ordinal irregularity and geometric roughness. These correlation patterns are qualitatively preserved across targets, supporting a "family-level" reading of subsequent permutation-importance rankings: rather than over-interpreting fine rank swaps among strongly correlated descriptors, we emphasise which descriptor families remain informative and stable across window choices (Section 3.1).

Finally, we stress that these correlation matrices are descriptive and not causal: even after differencing, residual autocorrelation and common drivers may induce nontrivial associations, and correlation alone does not identify physical mechanisms (Pearl, 2008, 2009). This motivates the combined use of (i) stability diagnostics across window lengths and (ii) model-agnostic attribution as complementary evidence, rather than relying on correlations as mechanistic proof.

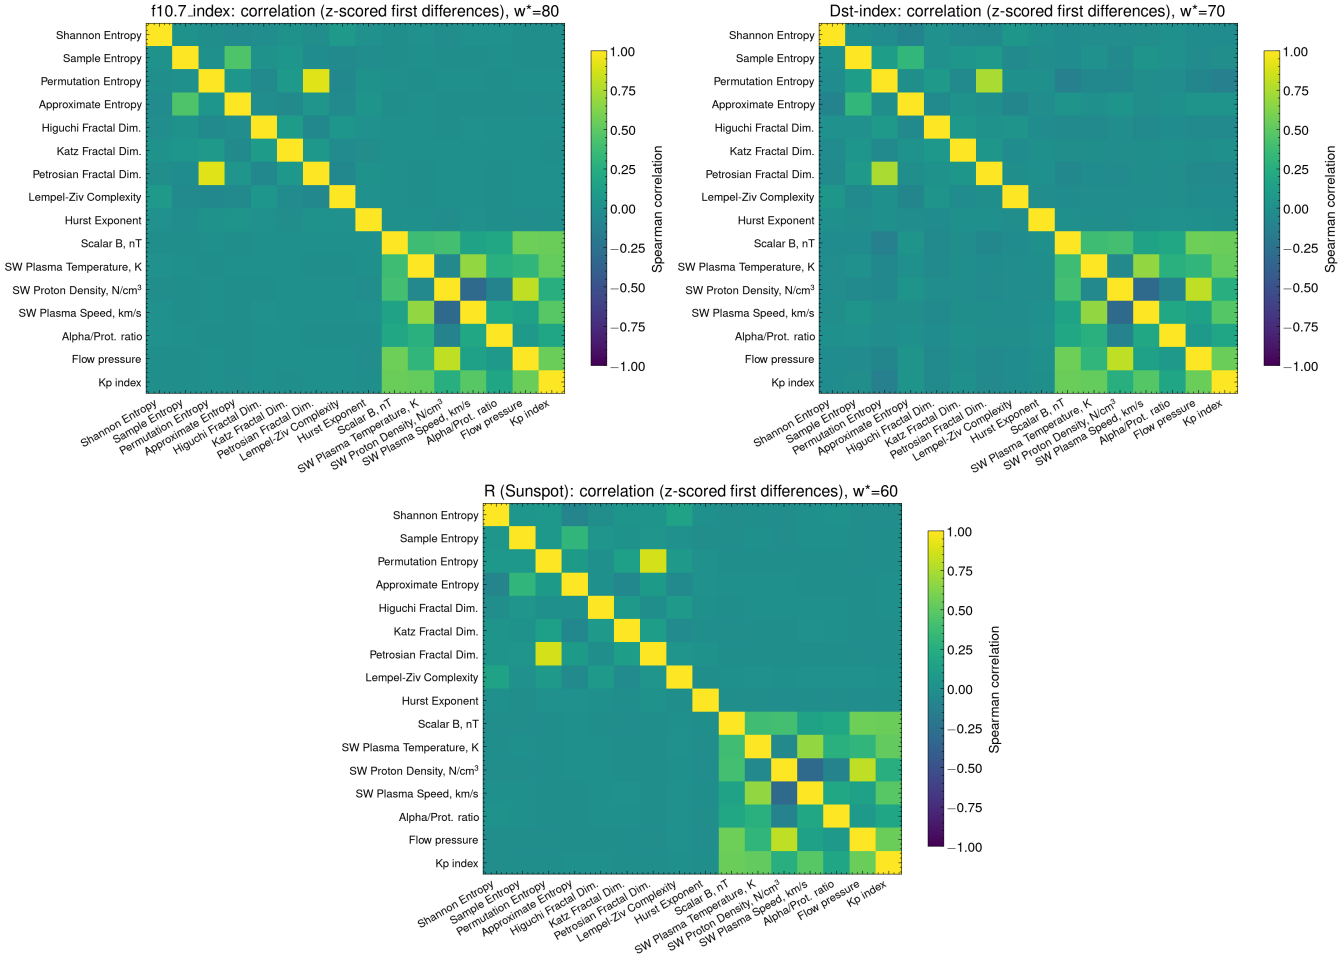


Fig. 2: Spearman correlation matrices computed on *standardised first differences* of the windowed features, using the target-specific windows $w^*_{\mathrm{F10.7}} = 80$ days (top-left), $w^*_{\mathrm{Dst}} = 70$ days (top-right) and $w^*_{\mathrm{SN}} = 60$ days (bottom-middle). This differenced view mitigates spurious trend-driven associations (Yule, 1926) and highlights the dominant redundancy structure: an entropy–roughness block (entropies and fractal dimensions) and a more weakly coupled pair (Lempel–Ziv complexity and Hurst exponent).

### 3.3. *Robustness of permutation-importance rankings across window lengths*

Using the target-specific design matrices (with the target amplitude excluded as a predictor), we compute permutation importance for both learners under a chronological train/test split and training-only preprocessing (Section 2.2.3). To mitigate instability driven by the sliding-window length, we repeat the full attribution analysis across the window grid used in Section 3.1 and quantify

robustness by counting how often each predictor appears among the top-10 features (across windows × models; maximum = 22 appearances). Figure 3 summarises these robustness counts for Dst, F10.7, and Sunspot Number.

Two consistent messages emerge. First, a compact set of predictors is remarkably stable across temporal context choices: several variables enter the top-10 almost systematically, indicating that the attribution results are not an artefact of a single window length. Second, the stable predictors differ by target in a physically interpretable way, suggesting that the learning models behave as diagnostic probes of distinct coupling pathways rather than as arbitrary function approximators. Because permutation importance can redistribute attribution among correlated predictors, we emphasise cross-window and cross-learner stability (Fig. 3) rather than over-interpreting small rank swaps within correlated families (Gneiting & Schlather, 2004).

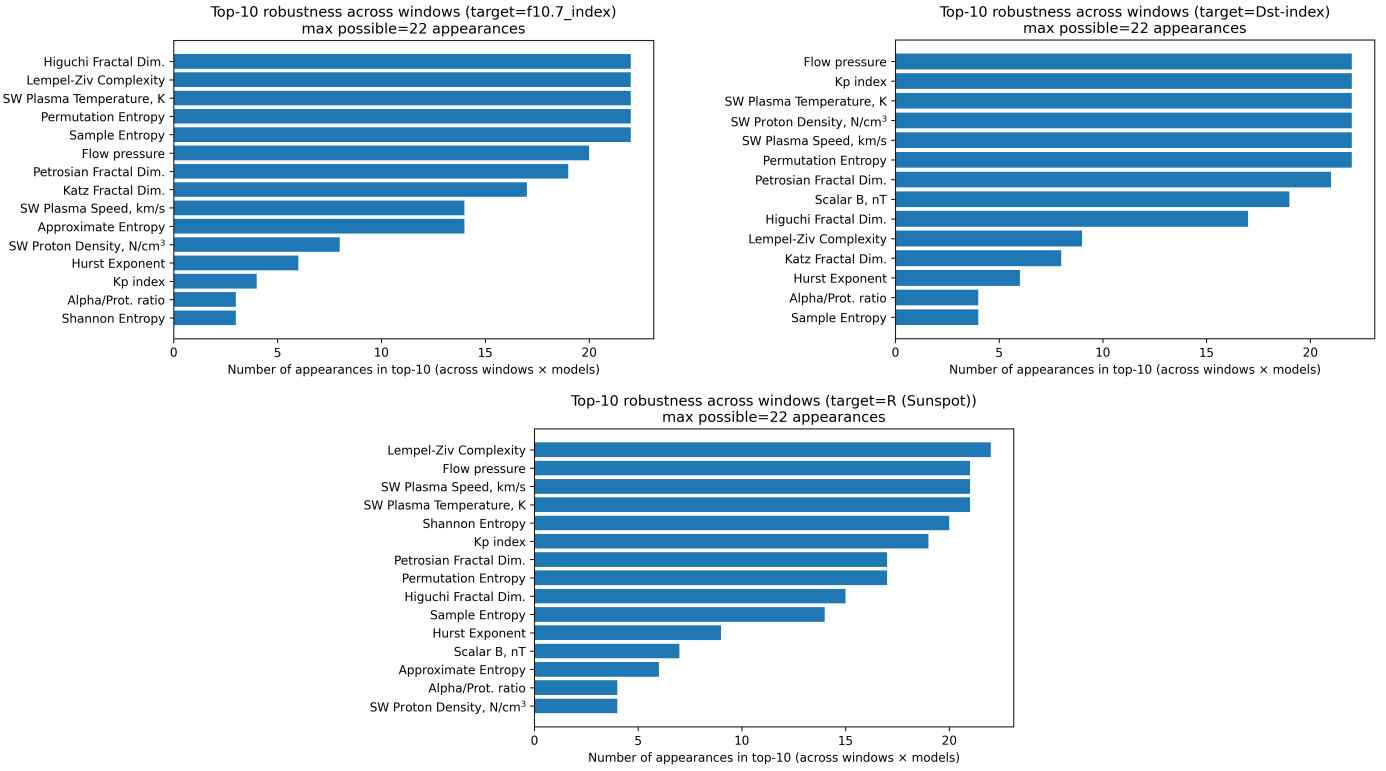


Fig. 3: Robustness of permutation-importance rankings across window lengths and learners, measured as the number of appearances of each predictor in the top-10 list (maximum = 22 appearances across windows × models), for three variables in panels: top-left: F10.7; top-right: Dst; bottom-middle: Sunspot Number. Persistent selection across panels indicates stability of attribution to the choice of temporal context, supporting family-level interpretation of the leading predictors.

For Dst, the most robust predictors are dominated by classical geoeffective drivers and geomagnetic context: flow pressure and the Kp index appear essentially maximally often, accompanied by solar-wind plasma temperature, density, and speed. This pattern is expected because Dst variability is tightly tied to solar-wind forcing and magnetospheric state. Importantly, however, at least one shape descriptor—Permutation entropy—is also maximally robust, and Petrosian/Higuchi-type roughness measures remain frequent. This indicates that, even when classical forcing dominates, the morphology of the Dst trajectory (irregularity/roughness within the past-only window) carries additional explanatory signal that is not fully absorbed by bulk plasma parameters.

For F10.7, the robust set shifts toward descriptors that track multi-scale roughness and pattern novelty: Higuchi fractal dimension and Lempel–Ziv complexity are among the most consistently selected predictors, together with permutation and sample entropy. The appearance of flow pressure and solar-wind temperature among the robust variables suggests that part of the F10.7 variability co-moves with heliospheric conditions, yet the repeated selection of roughness/complexity metrics implies that the radio-flux time series encodes regime-dependent texture beyond its amplitude. In other words, the models repeatedly exploit windowed shape information to account for changes in the organisation of fluctuations across cycles, not merely changes in mean level.

For Sunspot Number, Lempel–Ziv complexity is the single most stable descriptor, accompanied by Shannon entropy and multiple solar-wind covariates (flow pressure, speed, and temperature) as well as Kp. This mixed profile is consistent with sunspot variability reflecting both intrinsic solar-cycle modulation (captured by information-content/novelty metrics) and concurrent heliospheric conditions that co-vary with solar activity. Notably, the repeated presence of Petrosian/Higuchi dimensions and permutation entropy supports the interpretation that solar-activity proxies are not only large when active but also exhibit systematically different morphological regimes (burstiness, roughness, ordinal disorder) across phases of the cycle.

Finally, Fig. 3 supports a family-level conclusion: across all three targets, at least one complexity/entropy/roughness descriptor is robustly relevant, often alongside a small subset of classical covariates. This cross-target consensus motivates focusing subsequent interpretation on stable feature families (novelty/complexity, ordinal irregularity, and roughness) rather than on any single metric in isolation.

### *3.4. Temporal context: target series and leading shape signatures*

To visualise the dynamical meaning of the highest-ranking descriptors, Figure 4 juxtaposes each target with the time-aligned evolution of three leading shape signatures computed at the target-specific window length $w^*$. To avoid overplotting and axis crowding, we show one full solar cycle (Solar Cycle 24, 2009–2019) in the main text and provide the full 1964–2025 panels in the Supplementary Material. All series are standardised (z-scored) to emphasise co-variability in timing and morphology rather than absolute amplitude.

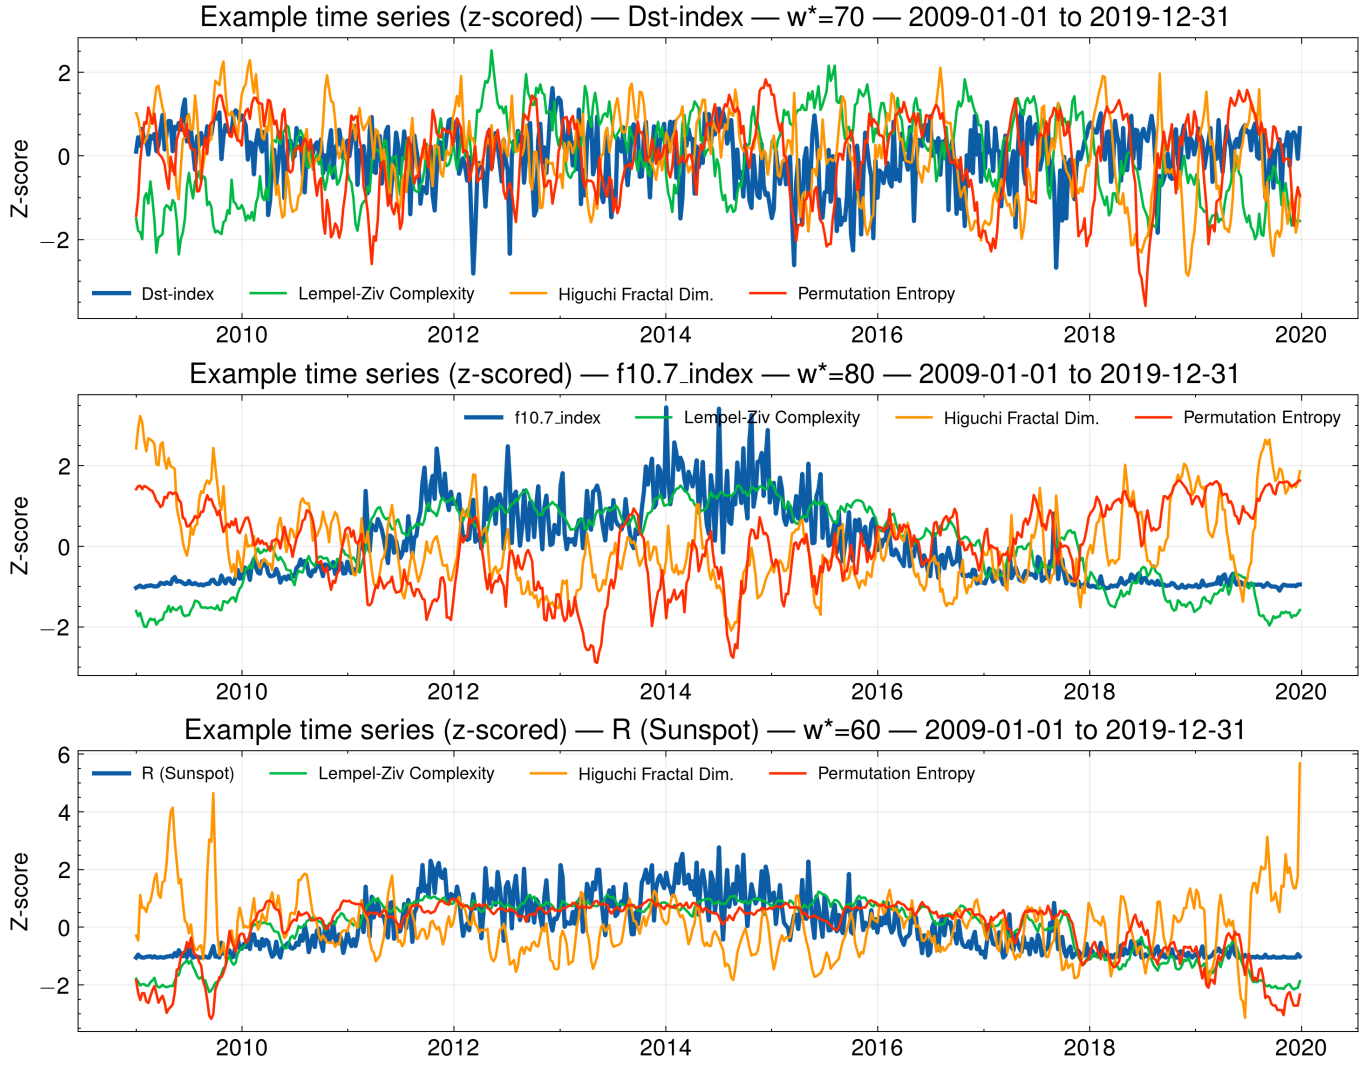


Fig. 4: Example time series (z-scored) for Solar Cycle 24 (2009–2019). Each panel overlays the target series with three leading shape descriptors computed using the target-specific window length $w^*$: Lempel–Ziv complexity, Higuchi fractal dimension, and permutation entropy. The trajectories are shown as descriptive dynamical fingerprints (not causal drivers) and illustrate that the dominant descriptors can be phase-shifted relative to the cycle-scale envelope while still tracking characteristic changes in irregularity and roughness.

Across the three targets, the selected descriptors highlight complementary aspects of variability. Lempel–Ziv complexity captures changes in algorithmic novelty (compressibility) of the windowed dynamics, permutation entropy reflects ordinal disorder, and Higuchi fractal dimension tracks multi-scale roughness. The resulting trajectories show that these notions of shape can evolve with clear phase structure over the cycle while also exhibiting shorter-term excursions, indicating that the information captured by the descriptors is not reducible to a simple rescaling of the target amplitude.

Importantly, the descriptor trajectories are not constrained to be in phase with the canonical solar-cycle envelope. In the solar proxies (F10.7 and sunspot number), the leading shape signatures display extended intervals of coherent evolution alongside intervals where roughness/irregularity measures vary more rapidly than the target itself, consistent with the idea that the descriptors encode morphology beyond level changes. For Dst, which is more intermittent and event-dominated, the same descriptors accentuate burstiness and irregularity episodes superposed on slower background variability, providing a compact, visually interpretable link between the permutation-importance results and the underlying time-domain behaviour.

These temporal patterns are descriptive and do not imply causality; they should be interpreted as dynamical fingerprints that complement amplitude-based summaries and motivate a family-level interpretation of relevance when descriptors are partially redundant.

## 4. Discussions

The F10.7 cm solar radio flux index has become a cornerstone in space-weather modelling, serving as a standard input parameter for numerous predictive frameworks. Its prominence is comparable to other established indices such as Dst or Kp that are used for geomagnetic activity assessment, positioning F10.7 as a widely adopted indicator of solar-activity conditions in operational models (Winter & Ledbetter, 2015). The index's long-term availability and consistent measurement methodology have contributed to its widespread adoption across diverse modelling applications.

In ionospheric modelling specifically, F10.7 has maintained its status as a predominant solar proxy, including in recently developed empirical models. This persistent usage stems from F10.7's ability to capture solar-cycle modulations that track EUV-related forcing on the upper atmosphere, making it valuable for parameterising solar influences in complex atmospheric simulations (Goncharenko et al., 2021; Laštovička, 2021; Laštovička & Burešová, 2023).

Recent advances in machine-learning approaches to space-weather forecasting have begun to challenge the traditional reliance on F10.7. Studies implementing gradient-boosting regressors have reported that models trained with F10.7 (or Mg II indices, or both) may perform no better on some out-of-sample settings than those without these indicators (Bailey et al., 2021). This finding suggests that, depending on the target and the learning setup, sufficient information about recurrent solar-wind/geomagnetic variability may already be encoded in other covariates, or can be partially reconstructed from dynamical signatures of the target itself.

Our results add an important nuance to this discussion: the strongest predictors are frequently not bulk levels themselves, but windowed *shape descriptors* computed from the past history of the target at the empirically selected $w^*$. Across the three targets, descriptors related to irregularity, roughness, and temporal organisation consistently appear among the leading predictors. This should not be interpreted as evidence that complexity drives solar or geomagnetic activity; rather, these descriptors provide compact nonlinear summaries of recent dynamical morphology (burstiness, intermittency, and ordering) that can complement, and sometimes partially substitute for, classical physical covariates in a predictive representation.

A particularly instructive example is Lempel–Ziv (LZ) complexity for F10.7. Figure 5 shows a strong monotonic relationship between windowed LZ and the contemporaneous F10.7 level over the full record. This indicates that LZ partially behaves as a nonlinear proxy of cycle phase/amplitude: as activity increases, the local morphological diversity of fluctuations tends to increase as well, and LZ captures this change in symbolic/algorithmic novelty within the $w^*$-day window. This provides a concrete explanation for why LZ ranks highly even when we exclude the target amplitude as an explicit predictor: the recent *shape* of the series carries amplitude-related information without collapsing it to a single scalar level.

A key question is whether the prominence of Lempel–Ziv (LZ) complexity reflects genuinely dynamical information or is largely a proxy for signal magnitude. We report the empirical relationship between windowed LZ complexity and the raw target amplitude for all three indices at their optimal window lengths. For F10.7 and Sunspot Number, LZ exhibits a strong monotone association with the cycle envelope (high Pearson/Spearman correlations and a clear saturation pattern at large amplitudes), consistent with the interpretation of LZ as a practical estimator of algorithmic complexity that increases as the signal becomes richer in distinct patterns and excursions (Kaspar & Schuster, 1987; Brudno, 1978). In these two cases, the tight coupling indicates that part of the relevance assigned to LZ by permutation importance can legitimately arise from its ability to encode the long-term activity state, while still preserving sensitivity to within-cycle irregularity.

In contrast, for Dst the amplitude–LZ association is weak (near-zero to mildly negative correlations), indicating that LZ is not merely re-expressing storm magnitude but rather responds to morphological aspects of geomagnetic variability such as intermittency, burst clustering, and changes in predictability. This cross-target contrast strengthens the transferability claim: the same descriptor family remains relevant across indices, yet the meaning of the LZ signal is index-dependent. From a dynamical-systems viewpoint, this is consistent with the role of LZ as an estimator of trajectory complexity/topological entropy for the underlying generator observed through scalar measurements, with interpretability mediated by the observable and its embedding (Sauer et al., 1991). Practically, these results motivate reporting LZ–amplitude diagnostics alongside importance rankings, and caution against interpreting LZ dominance as uniformly non-amplitude across all targets.

At the same time, the strong association in Fig. 5 motivates a cautious interpretation of permutation-importance ranks for LZ. Because LZ is not orthogonal to the cycle envelope, part of its relevance can reflect shared low-frequency structure rather than an independent physical driver. This aligns with the correlation and redundancy analysis: correlated descriptor families can redistribute

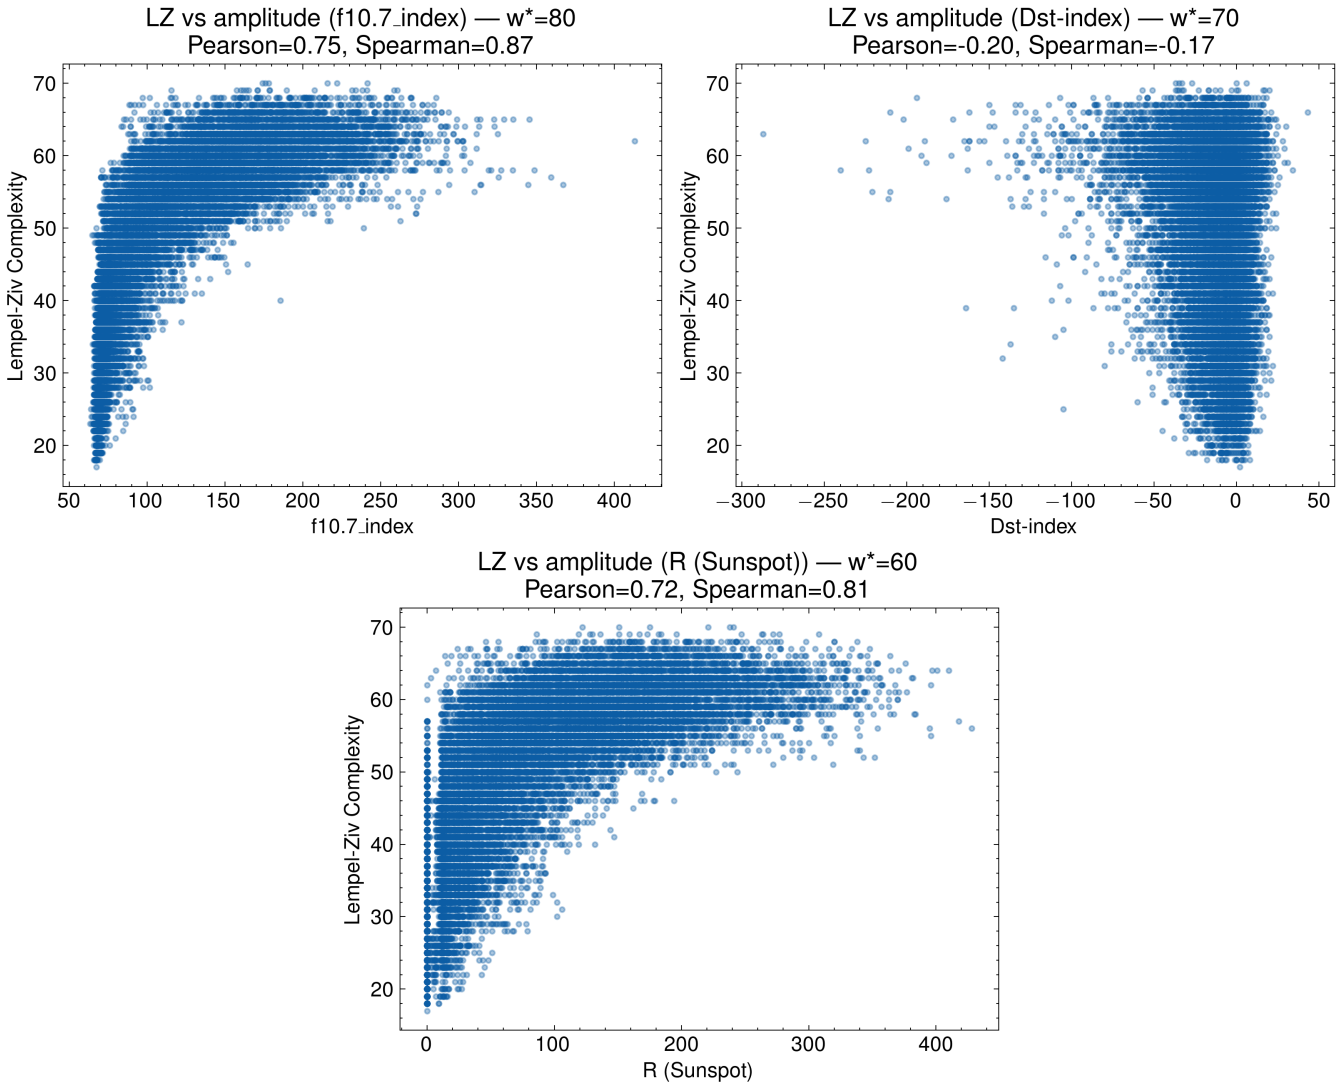


Fig. 5: Windowed Lempel–Ziv (LZ) complexity versus target amplitude for (top-left) F10.7 ($w^* = 80$), (top-right) Dst ($w^* = 70$) and (bottom-middle) Sunspot Number ($w^* = 60$). Pearson and Spearman correlations are reported in each panel. The strong monotone association for F10.7 and Sunspot contrasts with the weak dependence for Dst, indicating that LZ can act as an activity-state proxy for solar-cycle indices while remaining sensitive to morphology/intermittency for geomagnetic variability.

attribution under permutation, and therefore the most defensible statements are those that are stable across window lengths and learner classes, rather than small within-family rank swaps.

The preferred windows $w^*$ of ∼60–80 days suggest a natural heliophysical timescale for the representation: roughly two to three solar rotations. On this horizon, persistent active-region complexes and recurrent solar-wind structures can influence both solar proxies and geomagnetic indices, making windowed descriptors effective summaries of medium-term organisation without imposing a strict parametric cycle model. In this sense, our results support a representation in which robust predictive information is concentrated in a compact set of dynamical fingerprints that encode the recent temporal context, while avoiding causal claims from purely statistical associations.

From a physical perspective, the cycle-scale indices (F10.7 and Sunspot Number) encode the evolving population of magnetically active regions, for which both amplitude and temporal irregularity change systematically across the cycle. Entropy-, scaling-, and algorithmic-complexity descriptors provide a compact way to summarize intermittency, roughness, and pattern novelty induced by the emergence and decay statistics of active regions, beyond what a single amplitude level captures. For Dst, which is dominated by episodic storm-time dynamics, the weaker LZ–amplitude coupling suggests that complexity measures are sensitive to event clustering and temporal organization rather than storm magnitude alone.

Finally, we stress that correlation patterns between windowed descriptors and physical covariates should be interpreted cautiously. In oscillatory and trend-dominated geophysical records, correlations on levels may arise spuriously from shared nonstationarity or common cycle phases rather than from a direct physical linkage (Yule, 1926). Accordingly, we emphasize correlations

on standardized first differences and window-robust ranking stability as descriptive evidence, not as causal attribution. More generally, statistical association does not establish causal direction or mechanism, and any causal interpretation would require an explicit structural model and/or intervention-based identification (Pearl, 2009).

## 5. Conclusions

We presented a systematic assessment of windowed shape descriptors as diagnostic predictors of heliophysical variability across three canonical targets (F10.7, Sunspot Number, and Dst) over 1964–2025. Using a past-only sliding-window convention, we extracted a compact panel of morphology-oriented measures (entropies, fractal-dimension estimators, Lempel–Ziv complexity, and the Hurst exponent) and contrasted them with standard solar-wind covariates. Rather than proposing an operational forecaster, we used two complementary learners as attribution probes to obtain permutation-importance rankings under chronological splitting and training-only preprocessing. A sensitivity analysis showed that attribution stability and functional adequacy favour target-specific temporal contexts, with $w^*_{\mathrm{F10.7}} = 80$ days, $w^*_{\mathrm{SN}} = 60$ days, and $w^*_{\mathrm{Dst}} = 70$ days, and that the leading feature families remain qualitatively preserved across nearby windows.

Across targets, shape descriptors consistently appear among the most relevant predictors, often matching or exceeding the importance of classical bulk plasma parameters. The redundancy analysis further supports a family-level interpretation: several entropy–roughness descriptors form correlated blocks, so small rank swaps should not be over-interpreted, and the most defensible conclusions are those stable across window lengths and learner classes. Temporal visualisations over Solar Cycle 24 illustrate that the top-ranked descriptors act as dynamical fingerprints of recent organisation and intermittency, complementing amplitude-based summaries without implying causality. In particular, the strong monotonic association between Lempel–Ziv complexity and F10.7 level indicates that part of its relevance arises from encoding cycle phase/amplitude through nonlinear morphology, clarifying why it can rank highly even when the target amplitude is excluded as an explicit predictor.

These findings suggest that a compact, target-aware set of complexity and roughness measures can provide an informative representation of heliophysical time series that is not fully captured by traditional covariates alone. Future work will expand the descriptor basis with a small number of visibility-graph or network features, quantify incremental value under strict leakage controls, and test real-time feasibility for higher-cadence proxies and additional geomagnetic indices while maintaining the diagnostic, non-causal interpretation adopted here.

## Acknowledgements

M.C.V gratefully acknowledges the financial support received through the Programa Interinstitucional para el Fortalecimiento de la Investigación y el Posgrado del Pacífico (DELFÍN) under the 2025 call, which made his research stay at Universidad Tecnológica de Bolívar possible. The authors also wish to thank the UTB Dirección de Internacionalización y Cooperación for their logistical assistance and continuous guidance throughout the stay. Finally, D.S.P. thanks the UTB Research Department for their support and accompaniment throughout the entire investigation.

## References

Aguirre, L. A., & Letellier, C. (2011). Investigating observability properties from data in nonlinear dynamics. *Physical Review E—Statistical, Nonlinear, and Soft Matter Physics*, *83*(6), 066209. doi:https://doi.org/10.1103/PhysRevE.83.066209.

Ali, Z. A., Abduljabbar, Z. H., Tahir, H. A. et al. (2023). extreme gradient boosting algorithm with machine learning: A review. *Academic Journal of Nawroz University*, *12*(2), 320–334. doi:https://doi.org/10.25007/ajnu.v12n2a1612.

Altmann, A., Toloşi, L., Sander, O. et al. (2010). Permutation importance: a corrected feature importance measure. *Bioinformatics*, *26*(10), 1340–1347. doi:https://doi.org/10.1093/bioinformatics/btq134.

Bailey, R. L., Reiss, M., Arge, C. N. et al. (2021). Using gradient boosting regression to improve ambient solar wind model predictions. *Space Weather*, *19*(5), e2020SW002673. doi:https://doi.org/10.1029/2020SW002673.

Balasis, G., Daglis, I. A., Papadimitriou, C. et al. (2009). Investigating dynamical complexity in the magnetosphere using various entropy measures. *Journal of Geophysical Research: Space Physics*, *114*(A9). doi:https://doi.org/10.1029/2008JA014035.

Bandt, C., & Pompe, B. (2002). Permutation entropy: a natural complexity measure for time series. *Physical review letters*, *88*(17), 174102. doi:https://doi.org/10.1103/PhysRevLett.88.174102.

Braverman, V. (2015). Sliding window algorithms. In *Encyclopedia of algorithms* (pp. 1–6). Springer. doi:https://doi.org/10.1007/978-3-642-27848-8_797-1.

Brudno, A. A. (1978). The complexity of the trajectories of a dynamical system. *Russian Mathematical Surveys*, *33*(1), 197. URL: https://doi.org/10.1070/RM1978v033n01ABEH002243. doi:10.1070/RM1978v033n01ABEH002243.

Carbone, A., Castelli, G., & Stanley, H. E. (2004). Time-dependent hurst exponent in financial time series. *Physica A: Statistical Mechanics and its Applications*, *344*(1-2), 267–271. doi:https://doi.org/10.1016/j.physa.2004.06.130.

Carlsson, G., & Vejdemo-Johansson, M. (2021). *Topological data analysis with applications*. Cambridge University Press. URL: https://lccn.loc.gov/2021024970. doi:https://doi.org/10.1017/9781108975704.

Chazal, F., & Michel, B. (2021). An introduction to topological data analysis: fundamental and practical aspects for data scientists. *Frontiers in artificial intelligence*, *4*, 667963. doi:https://doi.org/10.3389/frai.2021.667963.

Chen, R. X. F. (2021). A brief introduction to shannon's information theory. URL: https://arxiv.org/abs/1612.09316. doi:https://doi.org/10.48550/arXiv.1612.09316. arXiv:1612.09316.

Chen, T., He, T., Benesty, M. et al. (2015). Xgboost: extreme gradient boosting. *R package version 0.4-2*, *1*(4), 1–4.

De Freitas, D., & De Medeiros, J. (2009). Nonextensivity in the solar magnetic activity during the increasing phase of solar cycle 23. *Europhysics Letters*, *88*(1), 19001. doi:https://doi.org/10.1209/0295-5075/88/19001.

Deshmukh, V., Baskar, S., Berger, T. et al. (2023). Comparing feature sets and machine-learning models for prediction of solar flares-topology, physics, and model complexity. *Astronomy & Astrophysics*, *674*, A159. doi:https://doi.org/10.1051/0004-6361/202245742.

Donner, R. V., Stolbova, V., Balasis, G. et al. (2018). Temporal organization of magnetospheric fluctuations unveiled by recurrence patterns in the dst index. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *28*(8). doi:https://doi.org/10.1063/1.5024792.

Fasy, B. T., Kim, J., Lecci, F. et al. (2014). Introduction to the r package tda. *arXiv preprint arXiv:1411.1830*, . doi:https://doi.org/10.48550/arXiv.1411.1830.

Gneiting, T., & Schlather, M. (2004). Stochastic models that separate fractal dimension and the hurst effect. *SIAM review*, *46*(2), 269–282. doi:https://doi.org/10.1137/S0036144501394387.

Goncharenko, L. P., Tamburri, C. A., Tobiska, W. K. et al. (2021). A new model for ionospheric total electron content: The impact of solar flux proxies and indices. *Journal of Geophysical Research: Space Physics*, *126*(2), e2020JA028466. doi:https://doi.org/10.1029/2020JA028466.

Gopinath, S., & Prince, P. (2017). Multifractal characteristics of magnetospheric dynamics and their relationship with sunspot cycle. *Advances in Space Research*, *59*(9), 2265–2278. doi:https://doi.org/10.1016/j.asr.2017.02.011.

Higuchi, T. (1988). Approach to an irregular time series on the basis of the fractal theory. *Physica D: Nonlinear Phenomena*, *31*(2), 277–283. doi:https://doi.org/10.1016/0167-2789(88)90081-4.

Kaspar, F., & Schuster, H. (1987). Easily calculable measure for the complexity of spatiotemporal patterns. *Physical review A*, *36*(2), 842. doi:https://doi.org/10.1103/PhysRevA.36.842.

Kilpua, E., Good, S., Ala-Lahti, M. et al. (2022). Structure and fluctuations of a slow icme sheath observed at 0.5 au by the parker solar probe. *Astronomy & Astrophysics*, *663*, A108. doi:https://doi.org/10.1051/0004-6361/202142191.

Kilpua, E. K., Good, S., Ala-Lahti, M. et al. (2024). Permutation entropy and complexity analysis of large-scale solar wind structures and streams. In *Annales Geophysicae* (pp. 163–177). Copernicus Publications Göttingen, Germany volume 42. doi:https://doi.org/10.5194/angeo-42-163-2024.

Kruse, R., Mostaghim, S., Borgelt, C. et al. (2022). Multi-layer perceptrons. In *Computational intelligence: a methodological introduction* (pp. 53–124). Springer. doi:https://doi.org/10.1007/978-3-030-42227-1_5.

Laštovička, J. (2021). The best solar activity proxy for long-term ionospheric investigations. *Advances in Space Research*, *68*(6), 2354–2360. doi:https://doi.org/10.1016/j.asr.2021.06.032.

Laštovička, J., & Burešová, D. (2023). Relationships between fof2 and various solar activity proxies. *Space Weather*, *21*(4), e2022SW003359. doi:https://doi.org/10.1029/2022SW003359.

Livadiotis, G. (2020). Nonextensive statistical mechanics: Equivalence between dual entropy and dual probabilities. *Entropy*, *22*(6), 594. doi:https://doi.org/10.3390/e22060594.

Maria, C., Boissonnat, J.-D., Glisse, M. et al. (2014). The gudhi library: Simplicial complexes and persistent homology. In *Mathematical Software–ICMS 2014: 4th International Congress, Seoul, South Korea, August 5-9, 2014. Proceedings 4* (pp. 167–174). Springer. doi:https://doi.org/10.1007/978-3-662-44199-2_28.

Mielniczuk, J., & Wojdyłło, P. (2007). Estimation of hurst exponent revisited. *Computational statistics & data analysis*, *51*(9), 4510–4525. doi:https://doi.org/10.1016/j.csda.2006.07.033.

Mumford, D. (2002). Pattern theory: the mathematics of perception. URL: https://arxiv.org/abs/math/0212400. arXiv:math/0212400.

Pavlos, E. G., Malandraki, O. E., Khabarova, O. V. et al. (2019). Non-extensive statistical analysis of energetic particle flux enhancements caused by the interplanetary coronal mass ejection-heliospheric current sheet interaction. *Entropy*, *21*(7), 648. doi:https://doi.org/10.3390/e21070648.

Pearl, J. (2008). Causal inference. In *Proceedings of the 2008th International Conference on Causality: Objectives and Assessment - Volume 6* COA'08 (p. 39–58). JMLR.org.

Pearl, J. (2009). *Causality*. Cambridge university press.

Perea, J. A., & Harer, J. (2015). Sliding windows and persistence: An application of topological methods to signal analysis. *Foundations of computational mathematics*, *15*, 799–838. doi:https://doi.org/10.1007/s10208-014-9206-z.

Pérez, J. B., Hauke, S., Lupo, U. et al. (2021). giotto-ph: a python library for high-performance computation of persistent homology of vietoris-rips filtrations. *arXiv preprint arXiv:2107.05412*, . doi:https://doi.org/10.48550/arXiv.2107.05412.

Petrosian, A. (1995). Kolmogorov complexity of finite sequences and recognition of different preictal eeg patterns. In *Proceedings eighth IEEE symposium on computer-based medical systems* (pp. 212–217). IEEE. doi:https://doi.org/10.1109/CBMS.1995.465426.

Popescu, M.-C., Balas, V. E., Perescu-Popescu, L. et al. (2009). Multilayer perceptron and neural networks. *WSEAS Transactions on Circuits and Systems*, *8*(7), 579–588.

Richman, J. S., Lake, D. E., & Moorman, J. R. (2004). Sample entropy. In *Methods in enzymology* (pp. 172–184). Elsevier volume 384. doi:https://doi.org/10.1016/S0076-6879(04)84011-4.

Richman, J. S., & Moorman, J. R. (2000). Physiological time-series analysis using approximate entropy and sample entropy. *American journal of physiology-heart and circulatory physiology*, *278*(6), H2039–H2049. doi:https://doi.org/10.1152/ajpheart.2000.278.6.H2039.

Sarlis, N., Livadiotis, G., McComas, D. et al. (2024). Persistent behavior in solar energetic particle time series. *The Astrophysical Journal*, *969*(1), 64. doi:https://doi.org/10.3847/1538-4357/ad479d.

Sauer, T., Yorke, J. A., & Casdagli, M. (1991). Embedology. *Journal of statistical Physics*, *65*(3), 579–616. doi:https://doi.org/10.1007/BF01053745.

Shannon, C. E. (1948). A mathematical theory of communication. *The Bell system technical journal*, *27*(3), 379–423. doi:https://doi.org/10.1002/j.1538-7305.1948.tb01338.x.

Sierra-Porta, D. (2022). On the fractal properties of cosmic rays and sun dynamics cross-correlations. *Astrophysics and Space Science*, *367*(12), 116. doi:https://doi.org/10.1007/s10509-022-04151-5.

Sierra-Porta, D. (2024). Relationship between magnetic rigidity cutoff and chaotic behavior in cosmic ray time series using visibility graph and network analysis techniques. *Chaos: An Interdisciplinary Journal of Nonlinear Science*, *34*(2). doi:https://doi.org/10.1063/5.0167156.

Sierra-Porta, D., Canedo Verdugo, M., & Herrera Acevedo, D. D. (2025). Global complexity signatures of solar cycles: A unified Entropy–Fractal survey of OMNI solar wind data (1964–2025). Submitted to Advances in Space Research.

Sierra-Porta, D., & Domínguez-Monterroza, A.-R. (2022). Linking cosmic ray intensities to cutoff rigidity through multifractal detrented fluctuation analysis. *Physica A: Statistical Mechanics and its Applications*, *607*, 128159. doi:https://doi.org/10.1016/j.physa.2022.128159.

Stumpo, M., Consolini, G., Alberti, T. et al. (2020). Measuring information coupling between the solar wind and the magnetosphere–ionosphere system. *Entropy*, *22*(3), 276. doi:https://doi.org/10.3390/e22030276.

Wawrzaszek, A., Modzelewska, R., Krasińska, A. et al. (2022). Fractal dimension analysis of earth magnetic field during 26 august 2018 geomagnetic storm. *Entropy*, *24*(5), 699. doi:https://doi.org/10.3390/e24050699.

Winter, L. M., & Ledbetter, K. (2015). Type ii and type iii radio bursts and their correlation with solar energetic proton events. *The Astrophysical Journal*, *809*(1), 105. doi:https://doi.org/10.1088/0004-637X/809/1/105.

Yule, G. U. (1926). Why do we sometimes get nonsense-correlations between time-series?–a study in sampling and the nature of time-series. *Journal of the royal statistical society*, *89*(1), 1–63. doi:https://doi.org/10.2307/2341482.

Zozor, S., Ravier, P., & Buttelli, O. (2005). On lempel–ziv complexity for multidimensional data analysis. *Physica A-statistical Mechanics and Its Applications*, *345*, 285–302. URL: https://api.semanticscholar.org/CorpusID:26659917. doi:https://doi.org/10.1016/J.PHYSA.2004.07.025.