

MORPHOLOGICAL FINGERPRINTS OF FORBUSH DECREASES AND THEIR RELATION TO GEOMAGNETIC STORM SEVERITY

JUAN D. PEREZ-NAVARRO^{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 and

* Corresponding author: perezjuan@utb.edu.co (Juan D. Perez-Navarro)

D. SIERRA-PORTA^{a,**} 

Universidad Tecnológica de Bolívar. Grupo de Investigación Gravitación y Matemática Aplicada - GIGMA & Grupo de Investigación Física Aplicada y Procesamiento de Imágenes y Señales - FAPIS. Parque Industrial y Tecnológico Carlos Vélez Pombo Km 1 Vía Turbaco. Cartagena de Indias, 130010, Colombia and

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

Version June 23, 2026

Abstract

Forbush decreases (FDs) are transient depressions in the galactic cosmic-ray flux observed by global neutron-monitor networks and are commonly associated with interplanetary disturbances driven by coronal mass ejections and related shocks. Despite extensive observational work, quantitatively comparing FD morphology across events and linking it to storm severity remains challenging due to heterogeneous station responses, coverage gaps, and the multivariate nature of the network. This work introduces a graph-based event representation in which each FD is mapped to an event network constructed from pairwise dissimilarities between station response time series. A controlled sparse backbone is obtained via the minimum spanning tree, enabling comparable event graphs across cases. From each graph, a compact set of geometric/topological fingerprints is computed, including global integration measures, spectral summaries, mesoscopic structure, centrality aggregates, and complexity descriptors.

Predictive skill is assessed using strict leave-one-event-out validation over a pre-defined grid of distance metrics and distance-domain transformations, with selection criteria fixed *a priori*. The proposed fingerprints exhibit measurable signal for three tasks: (i) multi-class classification of geomagnetic storm intensity (G3/G4/G5) with moderate but consistent performance and errors dominated by adjacent categories; (ii) stronger binary severity screening ($\geq G4$ vs. G3) with high sensitivity to severe events; and (iii) drop regression with partial least squares achieving positive explained variance relative to a fold-wise mean baseline. These results support the premise that FD event morphology leaves reproducible network signatures that can be captured by compact graph fingerprints, offering a complementary, interpretable framework for comparative FD characterization and for studying how heliospheric drivers imprint structure on global neutron-monitor responses.

Subject headings: Forbush decrease – space weather – neutron monitor – network graph

1. INTRODUCTION

Forbush decreases (FDs) are transient reductions in the galactic cosmic-ray (GCR) flux observed at Earth (Forbush 1937; Chilingarian *et al.* 2022). The physical mechanism behind Forbush decreases involves the magnetic field of plasma solar wind sweeping galactic cosmic rays away from Earth (Ghosh and Raychaudhuri 2007; Cane 2000). These events are associated with large-scale perturbations in the solar wind and interplanetary magnetic field, particularly non-recurrent decreases that are caused by transient interplanetary events related to coronal mass ejections (CMEs) from the Sun and their interplanetary counterparts (ICMEs) (Blanco *et al.* 2013; Huang *et al.* 2017).

It is worth emphasizing that the physical relationship between FD amplitude and geomagnetic storm severity is indirect: while both phenomena are driven by CME/ICME transients, storm intensity is particularly sensitive to the southward component of the interplanetary magnetic field ($B_z < 0$), whereas FD amplitude depends primarily on the magnetic barrier structure of the ICME sheath and flux rope, which modulates cosmic-ray access independently of B_z polarity (NOAA Space Weather Prediction Center 2024; Cane 2000).

Because neutron monitors register the integrated response of secondary particles generated in the atmosphere, FD signatures are routinely tracked in multi-station networks, and they remain a standard observable for studying the coupling between interplanetary transients and near-Earth particle environments (Chilingarian *et al.* 2025; Bahena Bias and Villaseñor 2011;

Juan D. Perez-Navarro: perezjuan@utb.edu.co
D. Sierra-Porta: dporta@utb.edu.co

Papaioannou *et al.* 2013). Despite decades of observational work, robust event-level characterization across stations continues to face practical challenges: FD morphologies vary with event type, viewing geometry, background conditions, and data completeness, while station responses differ systematically across the global network. Recent work has further characterized FDs using global neutron monitor networks (Chilingarian *et al.* 2025), muon telescope arrays (Bahena Bias and Villaseñor 2011), anisotropy measurements and leader fraction analysis (Papaioannou *et al.* 2013; Potgieter 2008; Ruffolo *et al.* 2016), and precursor signatures such as loss-cone enhancements and pre-event flux variations.

A key driver of inter-station heterogeneity is geomagnetic shielding. Station cutoff rigidity determines the minimum particle rigidity required to reach a given detector and therefore modulates both FD amplitude and timing (Yushkov 2026). As a result, FD strength is not solely an intrinsic property of the interplanetary driver; it is also filtered by the spatial distribution of observing stations and their rigidity thresholds (Laitinen and Dalla 2021; Ihongo *et al.* 2019; Okany 2025; Sierra-Porta and Domínguez-Monterroza 2022; Sierra-Porta 2024). Beyond cutoff rigidity, the asymptotic viewing direction of each station — which determines the portion of the sky sampled by the detector — introduces additional inter-station variability in FD onset timing and amplitude, since stations whose asymptotic directions become magnetically connected to the disturbed interplanetary region may register a precursor signature earlier than others (Smart and Shea 2005).

This motivates analyses that explicitly combine multi-station measurements with station metadata rather than treating each time series in isolation. In practice, however, many studies focus on a limited number of stations, compress the network into summary curves, or rely on pairwise comparisons without a unifying representation that is directly comparable across events.

Graph-based representations provide a principled way to summarize multivariate observations (Freitas Silva *et al.* 2025, 2021). Nodes can represent stations and edges can encode inter-station similarity or dissimilarity, yielding an event-specific network that can be characterized by well-defined structural descriptors. In the context of FD observations, such representations offer a pathway to convert multi-station time series into compact event fingerprints that can be compared across drivers, storm intensities, and event magnitudes (Perez-Navarro and Sierra-Porta 2026). A central methodological issue is comparability: naive similarity thresholds often generate overly dense graphs that trivialize many network metrics, whereas overly aggressive pruning can distort the underlying organization (Gutierrez Gomez *et al.* 2017). A reproducible graph-based FD framework must therefore specify (i) the event window and missing-data treatment, (ii) the similarity/distance definition, and (iii) a graph-construction rule that yields interpretable and comparable connectivity across events.

This work introduces an event-level network construction for FDs observed in the Neutron Monitor Database (NMDB). For each event, stations are treated as nodes and edges are derived from pairwise dissimilarities between station response series within a common event window, with missing values handled through alternative,

explicitly defined strategies (coverage-based station filtering or iterative imputation). In addition, transformations and optional scalings are applied in the distance domain to control dynamic range and enhance cross-event comparability. For visualization and controlled backbone comparison, minimum spanning trees (MST) are used to provide sparse (Stam *et al.* 2014), connected representations that highlight the dominant inter-station couplings while keeping the edge count fixed across events.

The scientific objective is to test whether graph fingerprints capture systematic differences between FD events stratified by geomagnetic storm intensity and by FD magnitude (drop). Three complementary predictive tasks are considered: multi-class intensity classification (G3/G4/G5), binary severity screening ($\geq G4$ vs. G3), and drop regression. Throughout this work, storm intensity is classified using the NOAA geomagnetic storm scale (G-scale), where G1 denotes minor storms and G5 denotes extreme events (NOAA Space Weather Prediction Center 2024). Evaluation is performed under strict leave-one-event-out validation, and model selection is restricted to a pre-defined, theory-motivated grid of pipeline choices with selection criteria fixed *a priori*.

The contributions of this study are threefold. First, it provides a reproducible pipeline that maps multi-station FD time series into event graphs with explicit control of preprocessing, missing-data handling, and distance-domain operations. Second, it defines a compact set of graph descriptors spanning integration, spectral structure, mesoscopic organization, centrality aggregates, and complexity measures, enabling event-level analysis in small-sample regimes with leakage-free validation. Third, it provides evidence that graph fingerprints derived from NMDB carry information about storm severity and FD magnitude and reveal rigidity-conditioned organization in representative event backbones, supporting graph-based characterization as a complementary tool for FD and space-weather studies.

2. DATA

We use hourly neutron-monitor count-rate records from the NMDB, selecting the subset of stations available during each Forbush decrease event (Steigies 2023, 2016; Mavromichalaki *et al.* 2011). For each event, we extract a common analysis window around the disturbance and assemble a time-aligned matrix whose columns correspond to stations and whose rows correspond to measurement times. Station coverage varies across events due to data availability and quality control.

To characterize the magnitude of the event at each station, we compute the percentage reduction in the cosmic-ray count rate with respect to a pre-event reference level defined as the mean count rate over a quiet period of 24–48 hours immediately preceding the FD onset. This reference is initially derived from NMDB event metadata and subsequently verified by visual inspection of each individual time series to confirm the absence of prior disturbances.

Each FD event is assigned a storm intensity class on the NOAA G-scale, compiled from widely used FD catalogs and peer-reviewed studies, complemented with operational space-weather reports from NASA and GOES-based products where applicable. The complete list of events, including dates, storm class, observed FD drop,

and data sources, is provided in Appendix A. A full description of the neutron monitor stations used — including cutoff rigidity, geographic coordinates, altitude, and mean count rate — is given in Appendix B.

Finally, we collect station metadata from NMDB and associated station documentation, including cutoff rigidity, geographic coordinates (latitude and longitude), and altitude. These attributes are used both for descriptive analysis and to evaluate rigidity-conditioned network structure and node-level roles across events. In addition to cutoff rigidity, we note that the asymptotic viewing direction of each station introduces further variability in FD onset timing across the network; this geometric effect is discussed in the context of the results in Section 5.

3. METHODS

3.1. Event windows and missing-data treatments

For each Forbush decrease (FD) event, a common analysis window is extracted from multi-station NMDB count-rate records and arranged into a time-aligned matrix $X \in \mathbb{R}^{T \times N}$, where T is the number of time samples and N is the number of stations available for that event.

Count-rate records are used at hourly resolution, consistent with the temporal scale of Forbush decrease dynamics. Where stations provide only minute-resolution data, hourly rates are obtained by averaging, which both reduces statistical noise and allows the inclusion of additional stations that report only at hourly cadence. A complementary missing-data treatments is considered to assess robustness. (i) Coverage filtering: stations whose fraction of missing values within the event window exceeds a fixed threshold $\tau = 0.5$ are excluded, and subsequent computations operate on the retained station set. (ii) Imputation: missing samples in X are imputed using an iterative multivariate scheme with a k -nearest neighbours (k -NN, $k = 2$) regressor as the base estimator (Van Buuren and Groothuis-Oudshoorn 2011; Pedregosa et al. 2011), producing a complete matrix prior to distance estimation. These two treatments define alternative event representations and are treated as part of the pre-defined pipeline grid.

3.2. Pairwise distances between station responses

Let $\mathbf{x}_i \in \mathbb{R}^T$ denote the response series for station i in a given event window. Each series $\mathbf{x}_i$ represents the percentage deviation of the hourly count rate from the station-specific pre-event reference level, as defined in Section 2. This representation places all stations on a common relative scale, making pairwise distances directly comparable across stations with very different absolute count rates (which depend on detector size, altitude, and tube configuration). Event-specific pairwise distances are computed between station responses to form a matrix $D \in \mathbb{R}^{N \times N}$ with entries d_{ij} . Two ℓ_p distances are considered:

$$d_{ij}^{(p)} = \left(\sum_{t=1}^T |x_i(t) - x_j(t)|^p \right)^{1/p}, \quad (1)$$

where $p = 1$ corresponds to Manhattan distance, and $p = 2$ corresponds to the Euclidean member of the Minkowski family (Shirkhorshidi et al. 2015; Wang et al. 2021; Shifaz et al. 2023). This choice provides a par-

simonious comparison between a more robust absolute-deviation geometry ($p = 1$) and a quadratic geometry ($p = 2$).

We note that this percentage-deviation representation does not fully equalize station responses: stations at higher cutoff rigidity exhibit systematically smaller percentage variations (both from FDs and from other modulation, with the notable exception of the diurnal variation, which is nearly rigidity-independent). An additional standardization step — e.g., per-station z -scoring of $\mathbf{x}_i$ prior to distance computation — could place all stations on a more equal footing and is a natural candidate for the pipeline grid in future work.

3.3. Distance-domain transformations and normalizations

To control dynamic range and improve comparability across events, additional operations are applied directly in the distance domain. Let t denote a robust scale defined as the median of the off-diagonal distances $\{d_{ij}\}_{i < j}$ for a given event. Three distance-domain transformations are explored:

$$\tilde{d}_{ij} = d_{ij} \quad (\text{none}), \quad (2)$$

$$\tilde{d}_{ij} = \log\left(\frac{d_{ij}}{t}\right) \quad (\log), \quad (3)$$

$$\tilde{d}_{ij} = \exp\left(\frac{d_{ij}}{t}\right) \quad (\text{exponential}). \quad (4)$$

After transformation, an optional normalization is applied to $\tilde{D}$ (e.g., none, min-max, z -score, robust scaling, or decimal scaling), yielding a final event-specific dissimilarity matrix used for graph construction.

Each transformation serves a distinct purpose. The identity (no transform) preserves absolute dissimilarity magnitudes, which may retain information about event amplitude. The log transform compresses the dynamic range of the distance matrix, reducing the influence of outlier station pairs and stabilizing topology- and centrality-based summaries across events with heterogeneous station coverage. The exponential transform amplifies large dissimilarities, sharpening contrasts between strongly and weakly coupled station pairs at the cost of increased sensitivity to outliers.

The optional post-transformation normalization can introduce systematic effects that interact with downstream graph metrics. Min-max scaling maps all events to the same distance range $[0, 1]$, removing cross-event amplitude information but improving comparability of graph topology. Z -score and robust scaling preserve relative structure while standardizing spread; decimal scaling reduces numerical magnitude without altering ratios. These effects are part of the explicitly modeled pipeline grid and are accounted for through the *a priori* model selection protocol described in Section 3.6.

3.4. Graph construction

Each event is represented as a weighted, undirected graph $G = (V, E, W)$ (Chung and Lu 2006; Kim and Wilhelm 2008) where nodes correspond to stations ($|V| = N$). Edges connect all station pairs (E is complete over the retained stations), and weights are given by the transformed/normalized dissimilarities ($w_{ij} = \tilde{d}_{ij}$). In this

representation, smaller weights indicate stronger similarity (shorter dissimilarity). No additional edge sparsification is imposed; comparability across events is primarily enforced through the station-level coverage filtering and the distance-domain scaling described above.

A consistent visual pattern is a core-periphery organization in which intermediate-rigidity stations tend to occupy bridging positions, while very low-rigidity stations (polar/high-latitude, darker colors) more often appear as peripheral branches. High-rigidity stations (warmer colors) show partial clustering and connect through a small number of intermediates, suggesting that similarity in rigidity can translate into similarity of FD response morphology and, consequently, shorter dissimilarities. This qualitative behavior motivates rigidity-aware summaries such as assortativity and centrality-rigidity relations at the event level.

Figure 1 provides an illustrative event-level network visualization for a representative FD. Nodes correspond to neutron-monitor stations and node color encodes the cutoff rigidity (GV). Edges correspond to the MST computed from the event-specific dissimilarity matrix, providing a sparse, connected backbone with fixed edge count ($N - 1$) and enabling direct visual comparison across distance metrics.

The two vertical panels highlight that changing the distance metric (Manhattan vs. Minkowski) can alter the backbone geometry (e.g., hubness and branch structure) while preserving the main large-scale station groupings.

For the illustrated event (drop $\approx 10.4\%$), the backbone already exhibits coherent connectivity across multiple rigidity regimes, but with visible heterogeneity in node roles: some stations act as connectors across subgroups, while others remain weakly coupled and attach as leaves. Across the full event set, the regression analysis formally tests whether such global organization patterns (e.g., efficiency, spectral complexity, centrality aggregation, and mixing structure) carry predictive signal for drop magnitude beyond a fold-wise baseline, rather than relying on individual visual examples.

For visualization and controlled backbone comparison, the minimum spanning tree (MST) is computed from the event-specific dissimilarity matrix using Kruskal’s algorithm (Kruskal 1956). The MST retains exactly $N - 1$ edges — the minimal connected subgraph that links all N stations while minimizing total edge weight. This produces a sparse, connected backbone that highlights the dominant inter-station dissimilarity structure while keeping the edge count fixed and independent of event-specific threshold choices. Cross-event comparability is thereby enforced by construction: all event graphs share the same topological constraint ($N - 1$ edges), and differences in fingerprints reflect genuine differences in inter-station organization rather than artifacts of density variation (Stam *et al.* 2014).

3.5. Event-level graph fingerprints

From each weighted event graph, a compact fingerprint vector is computed by aggregating node-level and spectral/mesoscopic properties into event-level descriptors. The reported feature set includes: global efficiency, Estrada index, modularity, assortativity, average Katz centrality, average closeness, average betweenness, and a Laplacian-based summary, together with additional

complexity descriptors: a Shannon entropy of the edge-weight distribution, a fractal-dimension estimate obtained via the Compact-Box-Burning (CBB) algorithm (Song *et al.* 2007) applied directly to the MST graph, varying the box-diameter threshold $\ell_b \in \{2, \dots, 11\}$, counting the minimum number of boxes required to cover all nodes at each scale, and extracting the slope of the $\log N_b$ vs. $\log \ell_b$ regression (so that $d_f = \text{slope}$); and a Hurst exponent estimated via rescaled-range (R/S) analysis (Schölzel 2019) applied to the ordered sequence of MST edge weights.

When centrality measures depend on path lengths, weighted shortest paths are computed using the dissimilarity weights. Event-level centralities are obtained by averaging node-wise values, producing comparable scalars across events. These descriptors are designed to capture complementary aspects of global integration (efficiency), spectral complexity (Estrada), mesoscopic structure (modularity), mixing patterns (assortativity), and heterogeneity/complexity (entropy/fractal/Hurst) in the station-response network. For a description of these descriptors, metrics and markers see for example Omar and Plapper (2020) and references therein.

3.6. Event labels, predictive tasks, and evaluation protocol

Each FD event is assigned (i) a storm intensity label on the NOAA G-scale and (ii) a continuous magnitude label given by the observed FD drop (%). Three tasks are considered: multi-class intensity classification (G3/G4/G5), binary severity screening ($\geq G4$ vs. G3), and drop regression.

Model selection is performed over a pre-defined, theory-motivated grid of pipeline choices (missing-data treatment, distance metric, distance-domain transformation, and optional normalization). Selection criteria are fixed *a priori*: macro-F1 for classification and R^2 for regression. Predictive evaluation is carried out with strict leave-one-event-out (LOEO) validation. For regression, performance is additionally compared against a fold-wise mean-predictor baseline computed on each training fold.

In LOEO, each of the n events is held out in turn as the test set while the remaining $n - 1$ events form the training set; this process is repeated n times and performance is aggregated across all held-out events. Any preprocessing steps that involve data-driven fitting — specifically, median imputation and feature standardization — are fitted exclusively on the training fold and applied to the held-out event, ensuring that no information from the test event leaks into the model. Given the modest sample size ($n \leq 34$), LOEO provides the least biased estimate of generalization error available without external validation data.

This evaluation design follows established practice for small-sample classification in space-physics contexts (Sierra-Porta 2024; Sierra-Porta and Domínguez-Monterroza 2022).

4. RESULTS

The model-selection grid spans distance-domain transformations, normalization choices, and two adjacency metrics, evaluated under strict LOEO with leakage-free preprocessing. Selection criteria were fixed *a priori* (macro-F1 for classification; R^2 for regression). Only the




FIG. 1.— Illustrative two event graph for 2012-03-08 (G3) and 2023-04-23 (G2) which are representative Forbush decreases. Nodes are stations colored by cutoff rigidity (GV). Edges correspond to the MST computed from the event-specific dissimilarity matrix, yielding a sparse, connected backbone with $N - 1$ edges that minimizes total distance. Left: Manhattan ($p = 1$); right: Minkowski/Euclidean ($p = 2$). While the main station groupings persist, the MST topology (branching patterns and connector nodes) changes with the distance metric.

best-performing configuration per task is reported here; the full sweep is available in the accompanying repository.

4.1. Rigidity-conditioned node roles in event graphs

To assess whether geomagnetic shielding shapes *local* station roles within event graphs, node-level descriptors were stratified by cutoff rigidity and summarized across events. Stations were grouped into three rigidity bands using fixed physical thresholds: low rigidity ($R_c < 3$ GV), high rigidity ($R_c \geq 6$ GV), and an intermediate group spanning $3 < R_c < 6$ GV. For each event, node metrics were aggregated within each band to obtain group-level summaries, and the resulting distributions across events were compared.

Figure 2 shows the across-event distributions for four representative descriptors (avg. katz, avg. closeness, avg. betweenness, and a Laplacian-based summary). A clear separation by rigidity emerges. Low-rigidity stations exhibit systematically larger avg. katz and avg. closeness across events, indicating that these stations tend to occupy more central positions in the event-network backbone and contribute disproportionately to global communicability-like structure. In the same group, avg. betweenness is also markedly higher, suggesting that low-rigidity stations more frequently act as bridges connecting subnetworks (or branches) in the backbone representation. Conversely, medium- and high-rigidity groups show substantially lower centrality levels, consistent with a more peripheral or locally confined role.

The Laplacian-based summary displays a complementary pattern: the intermediate-rigidity group ($3 \leq R_c < 6$ GV) exhibits larger values and broader variability than the low- and high-rigidity groups. This is physically interpretable in terms of the shielding structure of Earth’s magnetosphere: low-rigidity (polar) stations are most sensitive to GCR modulation and respond strongly to

any interplanetary barrier, making them dominant hubs in the dissimilarity backbone; intermediate-rigidity stations, partially shielded, mediate between the strongly modulated polar responses and the weakly modulated equatorial responses, and consequently tend to occupy spectral-bridging positions in the graph; high-rigidity stations are the least modulated and appear as peripheral nodes. This hierarchy — directly traceable to geomagnetic shielding — is precisely what makes rigidity-conditioned centrality a physically grounded fingerprint component. Notably, the two descriptors showing the strongest rigidity separation here (avg. Katz and avg. Laplacian) also rank among the top predictors in the intensity classification models (Table 5), directly linking physical interpretability to predictive performance.

4.2. Multi-class classification of storm intensity (G3/G4/G5)

The main classification task targets the categorical storm intensity associated with each Forbush event (G3/G4/G5). The best multi-class result was obtained using a log transformation in the distance domain, no post-distance normalization, and Minkowski-based adjacency, followed by linear discriminant analysis. Under LOEO ($n = 33$ events), the classifier achieved an accuracy of 0.576, balanced accuracy of 0.572, and macro-F1 of 0.575.

Table 1 reports the confusion matrix. Class-wise recalls are $\text{Rec}_{G3} = 6/10 = 0.60$, $\text{Rec}_{G4} = 8/13 = 0.62$, and $\text{Rec}_{G5} = 5/10 = 0.50$, indicating that performance is not driven by a single class. Misclassifications are dominated by neighboring categories (notably $G4 \leftrightarrow G5$), consistent with the expectation that storm severity forms a continuum and that graph fingerprints may partially overlap when intensities are close in magnitude. Figure 3 visualizes the same configuration in discriminant space (LD1–LD2), showing a moderate but nontrivial organization




FIG. 2.— Rigidity-conditioned distributions of node-role summaries across events. For each FD event graph, node-level metrics are aggregated within three cutoff-rigidity bands defined by fixed thresholds: low ($R_c < 3$ GV), medium ($3 \leq R_c < 6$ GV), and high ($R_c \geq 6$ GV). Boxplots summarize the distributions across events. Low-rigidity stations show systematically larger avg. katz, avg. closeness, and avg. betweenness, indicating more central and bridging roles in event backbones, while the Laplacian-based summary peaks for the medium-rigidity group.

by intensity.

TABLE 1
LOEO CONFUSION MATRIX FOR THE BEST MULTI-CLASS INTENSITY MODEL (LDA; G3/G4/G5; $n = 33$).

	Pred G3	Pred G4	Pred G5
True G3	6	1	3
True G4	1	8	4
True G5	2	3	5




FIG. 3.— LDA projection (LD1-LD2) for the best multi-class intensity pipeline (log distance transform, no normalization and Minkowski adjacency). Colors indicate storm classes (G3/G4/G5).

4.3. Binary severity screening: $\geq G4$ versus $G3$

A binary severity screening task was also considered, grouping strong storms (G4/G5) versus moderate storms (G3). The best-performing configuration for this task — log distance transform, decimal-scaling normalization,

and Minkowski adjacency — differs from the multiclass optimum, reflecting that the binary and multiclass tasks favor different representations of inter-station dissimilarity structure. Under LOEO ($n = 33$), LDA achieved an accuracy of 0.788, balanced accuracy of 0.735, and macro-F1 of 0.741 (Table 3).

Sensitivity to severe events is high: the true positive rate for $\geq G4$ is $20/23 = 0.87$. Specificity for G3 is lower ($5/10 = 0.50$), indicating that a subset of moderate events remains difficult to separate from the severe group. This asymmetry is consistent with the expectation that stronger disturbances imprint more distinctive global network signatures, while moderate events can occupy an intermediate regime in which fingerprints overlap with adjacent categories.

TABLE 2
LOEO CONFUSION MATRIX FOR BINARY SEVERITY SCREENING (LDA; $\geq G4$ vs. $G3$; $n = 33$).

	Pred G3	Pred $\geq G4$
True G3	5	5
True $\geq G4$	3	20

4.4. Drop regression with PLS (5 components)

The regression task evaluates whether graph fingerprints encode quantitative information about Forbush magnitude by predicting the observed drop (%). The best performance was obtained with Minkowski-based adjacency and no distance transformation or normalization, combined with partial least squares regression (PLS) using 5 components. Under LOEO ($n = 34$), the model achieved $R^2 = 0.350$, MAE = 2.41 percentage points, and RMSE = 3.06.

This outperformed a fold-wise mean-predictor baseline (LOEO), which yielded $R^2 = -0.062$, MAE = 2.69, and RMSE = 3.91. Figure 4 shows predicted versus observed drops, indicating measurable skill together with a typical regression-to-the-mean pattern: very large drops are underpredicted and small drops slightly overpredicted, which is expected given the limited sample size and the intrinsic variability linking event morphology to magnitude.

4.5. Summary of best configurations

Table 3 summarizes the best-performing configurations and LOEO performance. Overall, the results support the hypothesis that geometric/topological graph fingerprints extracted from multi-station Forbush observations contain reproducible signal for (i) stratifying storm intensity and (ii) partially explaining drop magnitude. Given the modest number of events, these findings are best interpreted as evidence of signal in the proposed fingerprints rather than as a final operational model; larger event sets will enable tighter uncertainty quantification and more definitive benchmarking.

4.6. Comparison with scalar baselines

To assess the practical advantage of the graph-based pipeline over simpler alternatives, we evaluated six scalar baselines under the same LOEO protocol and LDA classifier used for the graph fingerprints. Baselines include:




FIG. 4.— PLS (5 components) LOEO regression: predicted vs. observed Forbush drop (%) for the best pipeline (no transform; no normalization; Minkowski adjacency). Each point represents one held-out event. The dashed line indicates perfect agreement between predicted and observed values.

TABLE 3
BEST-PERFORMING CONFIGURATIONS (SELECTED BY LOEO PERFORMANCE) FOR EACH TASK.

Task	Best setting	LOEO performance
Multi-class intensity (G3/G4/G5)	LDA; log distance transform; no normalization; Minkowski adjacency	Acc = 0.576; Balanced Acc = 0.572; Macro F1 = 0.575 ($n = 33$)
Binary severity ($\geq G4$ vs G3)	LDA; log distance transform; Decimal-Scaling; Minkowski adjacency	Acc = 0.7878; Balanced Acc = 0.7347; MacroF1 = 0.7413 ($n = 33$)
Drop regression	PLS-5; no transform; no normalization; Minkowski adjacency	$R^2 = 0.350$; MAE = 2.41; RMSE = 3.06 ($n = 34$)

the maximum FD drop observed across polar stations ($R_c < 3$ GV), the minimum southward IMF B_z component, the mean interplanetary magnetic field magnitude $|\mathbf{B}|$, and the mean solar-wind speed, each extracted from the OMNI hourly dataset over a window spanning 6 hours before to 24 hours after the event onset — a duration consistent with the typical main-phase and early-recovery timescale of a geomagnetic storm; a two-dimensional feature vector formed by the joint pair (B_z minimum, $|\mathbf{B}|$ mean) over the same window, used directly as input to the LDA classifier; and the minimum Dst index. OMNI data were obtained from the GSFC/SPDF OMNIWeb interface (King and Papitashvili 2005).

Table 4 summarizes the results. All scalar baselines fall below the graph fingerprints across the three tasks. In the multiclass setting, the best scalar/vector baseline is the joint (B_z , $|\mathbf{B}|$) pair (macro-F1 = 0.561), which lies 0.014 points below the graph fingerprints (macro-F1 = 0.575); the minimum southward B_z alone reaches macro-F1 = 0.501. In the binary screening task, the same joint baseline reaches macro-F1 = 0.618 against 0.741 for the graph pipeline (+0.123).

The polar max-drop — the most direct simple measure of FD amplitude — performs below the stratified-random baseline in the multiclass task (macro-F1 = 0.285 < 0.333), confirming that raw amplitude at polar stations does not stratify storm intensity. For drop regression, all scalar baselines yield $R^2 < 0$, indicating they per-

form worse than a fold-wise mean predictor, whereas the graph PLS achieves $R^2 = 0.350$. In absolute terms, the graph PLS achieves the only positive R^2 among all methods; while its MAE (2.41) is comparable to or better than most baselines, the polar max-drop baseline shows a similarly low MAE (2.28) despite its negative R^2 (-0.007), illustrating that low absolute error alone does not imply predictive skill above a trivial mean-predictor.

These results demonstrate that the geometric and topological structure encoded by the graph fingerprints carries information about storm severity and FD magnitude that is not captured by any individual scalar measure derived from either cosmic-ray amplitude or interplanetary conditions.

TABLE 4

LOEO PERFORMANCE OF SCALAR/VECTOR BASELINES VERSUS GRAPH FINGERPRINTS. MC: MULTICLASS G3/G4/G5 (MACRO-F1); BIN: BINARY $\geq G4$ vs G3 (MACRO-F1); REG: DROP REGRESSION. ALL CLASSIFIERS USE LDA; REGRESSION USES LINEAR REGRESSION FOR BASELINES AND PLS-5 FOR GRAPH FINGERPRINTS. B_z , $|\mathbf{B}|$, AND SOLAR-WIND SPEED ARE COMPUTED OVER A -6H/+24H WINDOW CENTERED ON THE EVENT ONSET. RANDOM: STRATIFIED-RANDOM BASELINE. MAE AND RMSE IN PERCENTAGE POINTS.

Method	MC F1	Bin F1	Reg R^2	MAE	RMSE
Random (stratified)	0.333	0.500	—	—	—
Polar max-drop ($R_c < 3$ GV)	0.285	0.411	-0.007	2.28	3.80
B_z southward min (OMNI)	0.501	0.505	-0.061	2.70	3.91
$ \mathbf{B} $ mean (OMNI)	0.470	0.411	-0.070	2.76	3.92
SW speed mean (OMNI)	0.150	0.411	-0.006	2.38	3.80
(B_z , $ \mathbf{B} $) joint (2-feat)	0.561	0.618	-0.285	2.96	4.30
Dst index	0.402	0.377	-0.057	2.73	3.90
Graph fingerprints	0.575	0.741	0.350	2.41	3.06

4.7. Which graph descriptors drive the predictions?

To improve interpretability, descriptor relevance was estimated for the best-performing models by aggregating standardized coefficients across LOEO folds (Table 5). For multi-class intensity classification (G3/G4/G5), the dominant contributors were average Katz centrality (avg_katz), followed by the Estrada index, Laplacian-based summaries (avg_laplacian), entropy, and modularity. This ranking suggests that severity stratification is primarily supported by a combination of communicability-like/global connectivity structure (avg_katz), spectral organization (Estrada and Laplacian summaries), complexity/heterogeneity (entropy), and mesoscopic organization (modularity).

For the binary severity screening task ($\geq G4$ vs. G3), the same core set remains dominant (Table 5), with avg_katz again ranking first, followed by avg_laplacian, entropy, Estrada index, and modularity. The improved separation observed in the binary setting is therefore consistent with stronger shifts in global connectivity/communicability and distance-geometry/spectral structure between moderate and severe events, complemented by complexity and community organization.

For drop regression (PLS-5), the largest standardized coefficients were associated with assortativity (positive) and average betweenness (negative), followed by smaller

TABLE 5

MOST INFLUENTIAL GRAPH DESCRIPTORS FOR THE BEST-PERFORMING MODELS, ESTIMATED BY AGGREGATING STANDARDIZED COEFFICIENTS ACROSS LOEO FOLDS. FOR CLASSIFICATION, VALUES REPORT MEAN $\pm$ STD RELEVANCE SCORES; FOR REGRESSION, COEFFICIENTS ARE REPORTED AS MEAN $\pm$ STD (STANDARDIZED).

Task	Top descriptors (ranked)
G3/G4/G5 (LDA)	avg_katz (7.23 $\pm$ 1.05), estrada_index (4.98 $\pm$ 0.62), avg_laplacian (4.56 $\pm$ 0.57), entropy (4.04 $\pm$ 0.54), modularity (2.76 $\pm$ 0.48)
$\geq$ G4 vs G3 (LDA)	avg_katz (20.21 $\pm$ 3.64), avg_laplacian (12.23 $\pm$ 1.48), entropy (10.77 $\pm$ 1.36), estrada_index (8.08 $\pm$ 1.05), modularity (6.38 $\pm$ 0.90)
Drop (PLS-5)	assortativity (+4.41 $\pm$ 0.27), avg_betweenness (-4.34 $\pm$ 0.25), modularity (-0.65 $\pm$ 0.09), fractal (-0.65 $\pm$ 0.10), avg_laplacian (-0.57 $\pm$ 0.09)

but consistent contributions from modularity, fractal-dimension estimates, and Laplacian summaries (mostly negative; Table 5). This pattern indicates that drop magnitude is more strongly linked to mixing and bridging structure (assortativity and betweenness), with additional modulation from global organization and complexity captured by mesoscopic and spectral/complexity descriptors.

5. DISCUSSION

Before interpreting the classification results, it is worth contextualizing the physical relationship between Forbush decrease amplitude and geomagnetic storm severity. Both phenomena are driven by the same interplanetary transients (CMEs and their associated sheaths), but they probe different physical properties. Storm intensity on the NOAA G-scale is primarily controlled by the southward component of the interplanetary magnetic field ($B_z < 0$), which drives efficient energy transfer into the magnetosphere through dayside reconnection (Cane 2000). FD amplitude, by contrast, depends on the magnetic barrier presented by the ICME sheath and flux rope, which modulates cosmic-ray access through gradient and curvature drifts and diffusion independently of B_z polarity (Richardson and Cane 2011; Smart and Shea 2005). Consequently, large FDs can occur without a strong geomagnetic storm and vice versa, and the correlation between FD amplitude and storm class is moderate at best. This indirect coupling is consistent with our baseline results (Section 4.6): the minimum southward B_z — the most direct storm driver — achieves macro-F1 = 0.501 in the multiclass task, while the polar max-drop achieves only 0.285. Even the joint (B_z , $|\mathbf{B}|$) pair, which reaches macro-F1 = 0.561, remains below the graph fingerprints (0.575) in the multiclass task and is more clearly outperformed in binary screening (0.618 vs. 0.741). This indicates that, while interplanetary drivers alone carry meaningful information, the morphological structure encoded by the graph fingerprints captures additional, complementary information not reducible to a small set of in-situ scalar measurements.

The results indicate that event-level graph fingerprints derived from multi-station neutron-monitor responses encode reproducible information about both geomagnetic-storm severity and Forbush magnitude. The moderate yet consistent performance in the G3/G4/G5 task, together with confusion dominated by

adjacent categories, supports a physically plausible picture in which storm intensity behaves as an ordered continuum rather than a set of sharply separated regimes (NOAA Space Weather Prediction Center 2024). In this context, the LDA projection reveals partial organization in discriminant space, while overlap remains expected due to shared heliospheric drivers and heterogeneous station responses modulated by geomagnetic shielding (Cane 2000; Richardson and Cane 2011; Smart and Shea 2005).

The stronger performance observed in the binary screening task ($\geq$ G4 vs. G3) suggests that the proposed fingerprints are particularly sensitive to the structural signatures of more intense disturbances. This behavior is consistent with the expectation that stronger CME/shock-driven transients can produce more coherent and global organization across stations, whereas moderate events may occupy an intermediate regime in which morphology overlaps with neighboring categories. The observed asymmetry between sensitivity and specificity further motivates reporting both multi-class stratification (for physical gradation) and binary screening (for robust detection of stronger events).

The regression results demonstrate measurable skill for predicting FD drop relative to a fold-wise mean baseline, indicating that global network organization carries quantitative information about event magnitude. The residual regression-to-the-mean pattern for extremes is consistent with limited sample size and with the intrinsically noisy mapping between response morphology and magnitude when station coverage, rigidity distribution, and event geometry vary across cases (Mishev *et al.* 2024). These considerations suggest that drop prediction is feasible but likely benefits from larger event sets and additional physically informed covariates. It is also worth noting that event selection conditioned on geomagnetic storm class introduces a sample bias: large FDs without an associated storm, or storms without a prominent FD, are excluded by construction. Future work should assess whether including such events — or adopting an FD-amplitude-first selection strategy — changes the fingerprint structure and predictive relationships observed here.

An additional point concerns representation choices in the distance domain. The configuration that maximized intensity classification favored a log transform, which is compatible with stabilizing topology- and centrality-based summaries by compressing the distance dynamic range and reducing the influence of outliers (Box and Cox 1964). In contrast, the best-performing regression configuration favored leaving distances untransformed, suggesting that absolute dissimilarity magnitudes can retain information relevant to drop that may be attenuated by aggressive scaling. This divergence reinforces that different scientific targets (categorical severity versus continuous magnitude) can prefer different representations, and that distance-domain operations should be treated as explicit modeling choices rather than purely technical preprocessing steps.

Finally, the MST backbone provides a controlled sparse, connected representation that improves comparability across events. This choice also mitigates well-known sensitivity of many network metrics to graph density when graphs are produced by naive thresholding.

At the same time, the absence of cycles by construction implies that some mesoscopic descriptors are effectively probed through global and spectral summaries rather than through dense clustering structure. This motivates targeted future tests with fixed-density augmentations (e.g., MST plus a small number of shortest non-tree edges) to quantify the incremental value of cyclic structure while preserving cross-event comparability (Tuminello *et al.* 2007).

We further note that the pre-event reference level used here, derived from NMDB event metadata and verified by visual inspection (Section 2), would not be straightforwardly available in real time: NMDB’s catalogued FD events have not been updated for nearly two years, so an operational implementation would require an independent, automatically computed reference level. An additional source of inter-station variability not explicitly modeled here is the asymptotic viewing direction of each detector, which determines the portion of the heliosphere sampled and can introduce systematic differences in FD onset timing between stations whose asymptotic directions become magnetically connected to the disturbed interplanetary region at different times (Smart and Shea 2005). Incorporating viewing geometry as a node attribute — alongside cutoff rigidity — represents a natural extension that could sharpen the physical interpretation of node roles in future work.

6. CONCLUSIONS AND OUTLOOK

A graph-based approach was developed to characterize Forbush decreases using event-level networks constructed from pairwise dissimilarities between station response series. A controlled sparse backbone (minimum spanning tree) yields comparable connected graphs across events, from which compact geometric/topological fingerprints are computed. Model selection was performed over a pre-defined, theory-motivated grid of pipeline choices, with selection criteria fixed *a priori* (macro-F1 for classification; R^2 for regression) and evaluation carried out under strict LOEO validation.

The main findings are: (i) multi-class intensity classification (G3/G4/G5) exhibits moderate but consistent skill, with errors dominated by neighboring categories, supporting a continuum-like interpretation of storm severity in fingerprint space; (ii) binary severity screening ($\geq G4$ vs. G3) yields stronger discrimination and high sensitivity to severe events, indicating potential utility for robust event triage; (iii) drop regression demonstrates measurable predictive skill relative to a fold-wise mean baseline, implying that global network organization encodes quantitative information about FD magnitude, albeit with regression-to-the-mean tendencies for extremes; and (iv) graph fingerprints outperform all scalar and vector baselines — including polar maximum, minimum southward B_z , interplanetary field magnitude, solar-wind speed, Dst, and the joint $(B_z, |\mathbf{B}|)$ pair — across the three tasks, with the margin being modest in the multiclass setting (0.575 vs. 0.561 for the best baseline) but substantial in binary screening (0.741 vs. 0.618) and regression, where all baselines yield $R^2 < 0$ while the graph PLS achieves $R^2 = 0.350$, confirming that the topological and geometric structure of the multi-station response encodes information not captured by any single physical index.

A practical consideration concerns the temporal scope of the method. In its current form, the pipeline requires the complete event window to compute the distance matrix and graph fingerprints, making it a retrospective characterization tool rather than an operational forecasting system. This is by design: the goal is comparative morphological analysis across events, not real-time prediction. Nevertheless, a progressive variant is conceivable — computing fingerprints on a rolling window as the FD develops and tracking their evolution — which could in principle provide an early severity assessment before the event concludes. Such an extension would require careful treatment of the window length and of the partial coverage problem (stations that come online or drop out mid-event), and is left as a direction for future work. We also note that existing real-time FD warning systems, such as that operated by the Australian Bureau of Meteorology Space Weather Services (Australian Bureau of Meteorology Space Weather Services 2025), rely on single-station or small-network triggers; the graph-based framework could complement these by providing richer morphological context once sufficient multi-station data are available.

Given the modest number of labeled events, these results are best interpreted as evidence of reproducible signal in the proposed fingerprints rather than as a final operational model. Several directions can strengthen both statistical confidence and physical interpretability: (a) **Larger and more diverse event sets:** expand the labeled catalog to stabilize model selection, enable calibration across solar-cycle phases, and support ordinal or cost-sensitive formulations aligned with the ordered nature of intensity classes; (b) **Uncertainty-aware benchmarking:** report confidence intervals via event-level resampling (bootstrap) and assess robustness under alternative splits (e.g., leave-one-year-out or leave-one-solar-cycle-out) when sample size permits; (c) **Controlled cycle reintroduction:** compare MST fingerprints against fixed-density alternatives (MST+ m edges, k -NN graphs) to quantify the incremental value of cyclic and mesoscopic structure while preserving comparability; (d) **Physics-guided hybridization:** combine graph fingerprints with contemporaneous heliospheric and geomagnetic covariates (e.g., Dst, solar-wind speed, dynamic pressure, IMF B_z) to evaluate incremental predictive value and improve attribution; (e) **Rigidity-aware interpretability:** incorporate station metadata (cutoff rigidity, latitude/longitude, altitude) as explicit model inputs or as conditioning variables to clarify how geomagnetic shielding shapes network roles and event fingerprints; (f) **Progressive event fingerprinting:** extend the pipeline to compute fingerprints on rolling windows during event development, enabling severity assessment before the event concludes and opening a pathway toward near-real-time application; and (g) **Viewing-geometry integration:** incorporate asymptotic viewing directions as node attributes alongside cutoff rigidity to sharpen the physical interpretation of node roles and account for onset-timing heterogeneity across the station network.

Overall, graph fingerprints provide a compact, interpretable representation of FD morphology that complements standard time-series and index-based analyses. With expanded catalogs, uncertainty quantification, and physics-informed integration, this framework can mature

into a practical tool for comparative FD characterization and for probing how heliospheric drivers imprint distinct network signatures on the global neutron-monitor response.

DATA AND CODE AVAILABILITY

Neutron-monitor count-rate data are publicly available through the NMDB database (www.nmdb.eu). Interplanetary magnetic field, solar-wind, and geomagnetic index data were obtained from the OMNI hourly dataset via the GSFC/SPDF OMNIWeb interface (<https://omniweb.gsfc.nasa.gov>). The analysis scripts, pro-

cessed feature tables, and baseline comparison notebook are available in the repository: <https://github.com/dosquisd/NMDB-FD-PredictorWithGraphs>

ACKNOWLEDGMENTS

We gratefully acknowledge the Dirección de Investigaciones at Universidad Tecnológica de Bolívar for their support and accompaniment throughout this research process. We acknowledge the NMDB database (www.nmdb.eu), founded under the European Union's FP7 programme (contract no. 213007), for providing data.

APPENDIX

EVENT CATALOG

Table 6 lists the 34 Forbush decrease events used in this study, ordered chronologically. For each event the date of the FD onset is given together with the associated NOAA geomagnetic storm class (G-scale), the observed network-level FD drop (%), and the minimum Dst index recorded during the event window (nT). Storm classes were compiled from the NOAA Space Weather Prediction Center storm archive, the Forbush Decrease and Interplanetary Events Database (FEID), and peer-reviewed catalogues; for events with conflicting classifications, the most conservative peer-reviewed source was adopted.

TABLE 6

FORBUSH DECREASE EVENTS USED IN THIS STUDY. DROP: NETWORK-LEVEL MAXIMUM PERCENTAGE REDUCTION IN COSMIC-RAY COUNT RATE. DST_{min} : MINIMUM HOURLY DST INDEX DURING THE EVENT WINDOW (OMNI/WDC).

Date	Class	Drop (%)	Dst_{min} (nT)
2001-04-08	G3	6.68	-100
2001-04-11	G4	12.71	-271
2001-04-28	G4	7.74	-130
2001-09-25	G4	8.13	-102
2001-10-11	G4	6.46	-46
2001-10-21	G5	6.94	-412
2001-11-06	G4	12.52	-292
2001-11-24	G5	9.25	-221
2002-04-17	G3	6.56	-126
2002-05-23	G3	6.93	-109
2002-11-16	G3	7.30	-106
2003-05-29	G4	6.45	-144
2003-10-29	G5	26.06	-383
2003-10-30	G5	8.67	-383
2003-11-20	G5	6.58	-422
2004-01-22	G3	8.71	-149
2004-07-26	G4	13.48	-200
2004-11-07	G5	8.11	-373
2004-11-09	G5	8.12	-383
2005-01-18	G3	6.19	-121
2005-01-21	G4	11.26	-101
2005-05-15	G5	11.45	-263
2005-08-24	G5	7.16	-216
2005-09-11	G3	12.25	-139
2006-12-14	G4	9.29	-146
2011-10-24	G3	6.01	-137
2012-03-08	G3	10.40	-148
2012-07-14	G4	6.70	-139
2015-06-22	G4	8.57	-198
2017-09-07	G4	6.41	-144
2021-11-03	G3	9.68	-117
2023-04-23	G2	6.57	-213
2024-03-24	G4	11.47	-130
2024-05-10	G5	15.32	-412

Class distribution: $G2 \times 1$, $G3 \times 10$, $G4 \times 13$, $G5 \times 10$.

NEUTRON MONITOR STATION CATALOG

Table 7 lists the 49 neutron monitor stations used across the 34 events. Station codes follow the NMDB convention. Not all stations contributed data to every event; the set of active stations varies with data availability and the coverage-

filtering criterion ($\tau = 0.5$). Cutoff rigidities (R_c) are vertical geomagnetic cutoff rigidities from the NMDB station documentation. Geographic coordinates and altitudes are taken from NMDB metadata at www.nmdb.eu.

TABLE 7

NEUTRON MONITOR STATIONS USED IN THIS STUDY, SORTED BY CUTOFF RIGIDITY. LAT/LON IN DECIMAL DEGREES (NEGATIVE = S/W). ALT: ALTITUDE ABOVE SEA LEVEL (M). R_c : VERTICAL GEOMAGNETIC CUTOFF RIGIDITY (GV). STATION CODES FOLLOW NMDB CONVENTIONS.

Code	Station / Location	Lat (°)	Lon (°)	Alt (m)	R_c (GV)
TERA	Terra Nova Bay, Antarctica	−74.69	164.12	10	0.01
DOMB	Dome B, Antarctica	−77.05	94.92	3650	0.01
DOMC	Dome C (Concordia), Antarctica	−75.10	123.35	3233	0.01
MRNY	Mirny, Antarctica	−66.55	93.02	35	0.03
SOPO	South Pole, Antarctica	−90.00	0.00	2820	0.10
SOPB	South Pole B, Antarctica	−90.00	0.00	2820	0.10
NEU3	Neumayer III, Antarctica	−70.67	−8.27	42	0.10
MWSB	Mawson B, Antarctica	−67.60	62.87	30	0.22
MWSN	Mawson, Antarctica	−67.60	62.87	30	0.22
FSMT	Fort Smith, Canada	60.02	−111.93	205	0.30
THUL	Thule, Greenland	76.50	−68.70	260	0.30
PWNC	Peawanuck, Canada	54.98	−85.43	53	0.30
NAIN	Nain, Canada	56.55	−61.68	46	0.30
JBGO	Jang Bogo, Antarctica	−74.62	164.22	29	0.30
INVK	Inuvik, Canada	68.35	−133.72	21	0.30
TXBY	Tixie Bay, Russia	71.58	128.92	0	0.48
NRLK	Norilsk, Russia	69.27	88.05	45	0.63
APTY	Apatity, Russia	67.57	33.40	177	0.65
OULU	Oulu, Finland	65.05	25.47	15	0.81
CALG	Calgary, Canada	51.08	−114.13	1128	1.08
KERG	Kerguelen Islands, France	−49.35	70.26	33	1.14
YKTK	Yakutsk, Russia	61.95	129.72	118	1.65
MGDN	Magadan, Russia	59.95	150.77	220	2.10
KIEL2	Kiel2, Germany	54.33	10.13	54	2.36
KIEL	Kiel, Germany	54.33	10.13	54	2.36
NEWK	Newark, USA	39.68	−75.75	50	2.40
MOSC	Moscow, Russia	55.47	37.32	200	2.43
MCRL	Mc Murdo, Antarctica	−77.85	166.72	48	2.46
DRBS	Dourbes, Belgium	50.10	4.60	225	3.18
IRK3	Irkutsk-3, Russia	52.27	104.32	435	3.64
IRK2	Irkutsk-2, Russia	52.27	104.32	435	3.64
IRKT	Irkutsk, Russia	52.27	104.32	435	3.64
LMKS	Lomnický Štít, Slovakia	49.20	20.22	2634	3.84
JUNG	Jungfrauoch, Switzerland	46.55	7.98	3475	4.49
JUNG1	Jungfrauoch-1, Switzerland	46.55	7.98	3475	4.49
HRMS	Hermanus, South Africa	−34.42	19.22	26	4.58
BKSN	Baksan, Russia	43.28	42.69	1700	5.70
AATB	Almaty B, Kazakhstan	43.25	76.92	800	5.90
AATA	Almaty A, Kazakhstan	43.25	76.92	800	5.90
ROME	Rome, Italy	41.86	12.47	60	6.27
CALM	Alma-Ata (historical), Kazakhstan	43.25	76.92	800	6.95
PTFM	Plateau de Bure, France	44.63	5.91	2552	6.98
ARNM	Arnhem, Netherlands	51.98	5.88	15	7.10
NANM	Naan, Netherlands	51.98	5.88	15	7.10
MXCO	Mexico City, Mexico	19.33	−99.18	2274	8.28
ATHN	Athens, Greece	37.97	23.72	260	8.53
TSMB	Tsumeb, Namibia	−19.20	17.58	1240	9.15
ESOI	Ein Shemer, Israel	32.50	35.00	70	10.75
DJON	Djourney, South Korea	—	—	—	11.20

[†] Coordinates and altitudes from NMDB documentation (www.nmdb.eu).

REFERENCES

- S. E. Forbush, *Physical Review* **51**, 1108 (1937).
A. Chilingarian, G. Hovsepyan, H. Martoyan, T. Karapetyan, B. Sargsyan, N. Nokolova, H. Angelov, D. Haas, J. Knapp, M. Walter, O. Ploc, J. Shlegl, M. Kakona, and I. Ambrosova, *arXiv e-prints*, [arXiv:2212.13514](https://arxiv.org/abs/2212.13514) (2022), [arXiv:2212.13514](https://arxiv.org/abs/2212.13514) [physics.space-ph].
K. Ghosh and P. Raychaudhuri, *arXiv e-prints*, [astro-ph/0701860](https://arxiv.org/abs/astro-ph/0701860) (2007), [arXiv:astro-ph/0701860](https://arxiv.org/abs/astro-ph/0701860) [astro-ph].
H. V. Cane, *Space Sci. Rev.* **93**, 55 (2000).
J. J. Blanco, E. Catalán, M. A. Hidalgo, J. Medina, O. García, and J. Rodríguez-Pacheco, *Sol. Phys.* **284**, 167 (2013), [arXiv:1302.2597](https://arxiv.org/abs/1302.2597) [astro-ph.SR].
D.-H. Huang, H.-Q. Hu, J.-L. Zhang, H. Lu, D.-L. Zhang, B.-S. Xue, and J.-T. Lu, *arXiv e-prints*, [arXiv:1710.00945](https://arxiv.org/abs/1710.00945) (2017), [arXiv:1710.00945](https://arxiv.org/abs/1710.00945) [astro-ph.HE].
NOAA Space Weather Prediction Center, “Noaa space weather scales (geomagnetic storms g1–g5),” (2024), accessed 2026-02-12.

- A. Chilingarian, T. Karapetyan, and B. Sargsyan, [arXiv e-prints](#), [arXiv:2506.17917 \(2025\)](#), [arXiv:2506.17917 \[astro-ph.SR\]](#).
- A. Bahena Bias and L. Villaseñor, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 287 (IOP, 2011) p. 012034.
- A. Papaioannou, A. Belov, H. Mavromichalaki, E. Eroshenko, V. Yanke, E. Asvestari, A. Abunin, and M. Abunina, in *Journal of Physics Conference Series*, Journal of Physics Conference Series, Vol. 409 (IOP, 2013) p. 012202.
- M. S. Potgieter, *Journal of Atmospheric and Solar-Terrestrial Physics* **70**, 207 (2008).
- D. Ruffolo, A. Sáiz, P.-S. Mangeard, N. Kamyran, P. Muangha, T. Nutaro, S. Sumran, C. Chaiwattana, N. Gasiprong, C. Channok, C. Wuttiya, M. Rujiwarodom, P. Tooprakai, B. Asavapibhop, J. W. Bieber, J. Clem, P. Evenson, and K. Munakata, *ApJ* **817**, 38 (2016).
- B. Y. Yushkov, *Advances in Space Research* **77**, 3549 (2026).
- T. Laitinen and S. Dalla, *ApJ* **906**, 9 (2021).
- G. Ihongo, D. Ruffolo, A. Saiz, U. Tortempun, and A. C. L. Chian, in *36th International Cosmic Ray Conference (ICRC2019)*, International Cosmic Ray Conference, Vol. 36 (2019) p. 1084.
- C. Okany, in *General Assembly and 5th Annual Conference of the African Astronomical Society* (2025) p. 27.
- D. Sierra-Porta and A.-R. Domínguez-Monteroza, *Physica A Statistical Mechanics and its Applications* **607**, 128159 (2022).
- D. Sierra-Porta, *Chaos* **34**, 023114 (2024).
- D. Smart and M. Shea, *Advances in Space Research* **36**, 2012 (2005).
- V. Freitas Silva, M. Eduarda Silva, P. Ribeiro, and F. Silva, *Data Mining and Knowledge Discovery* **39**, 17 (2025).
- V. Freitas Silva, M. Eduarda Silva, P. Ribeiro, and F. Silva, *Wiley Interdisciplinary Reviews: Data Mining and Knowledge Discovery* **11**, e1404 (2021).
- J. D. Perez-Navarro and D. Sierra-Porta, *Open Journal of Astrophysics* **9** (2026), 10.33232/001c.157585.
- L. Gutierrez Gomez, B. Chiem, and J.-C. Delvenne, [arXiv e-prints](#), [arXiv:1705.10817 \(2017\)](#), [arXiv:1705.10817 \[stat.ML\]](#).
- C. Stam, P. Tewarie, E. Van Dellen, E. Van Straaten, A. Hillebrand, and P. Van Mieghem, *International Journal of Psychophysiology* **92**, 129 (2014).
- C. T. Steigies, in *Data, Analysis and Software in Heliophysics (DASH) 2023* (2023) p. 19.
- C. T. Steigies, in *AGU Fall Meeting Abstracts*, AGU Fall Meeting Abstracts, Vol. 2016 (2016) pp. IN44A-05.
- H. Mavromichalaki, A. Papaioannou, C. Plainaki, C. Sarlanis, G. Souvatzoglou, M. Gerontidou, M. Papailiou, E. Eroshenko, A. Belov, V. Yanke, E. O. Flückiger, R. Bütikofer, M. Parisi, M. Storini, K.-L. Klein, N. Fuller, C. T. Steigies, O. M. Rother, B. Heber, R. F. Wimmer-Schweingruber, K. Kudela, I. Strharsky, R. Langer, I. Usoskin, A. Ibragimov, A. Chilingaryan, G. Hovsepyan, A. Reymers, A. Yeghikyan, O. Kryakunova, E. Dryn, N. Nikolayevskiy, L. Dorman, and L. Pustil'Nik, *Advances in Space Research* **47**, 2210 (2011).
- S. Van Buuren and K. Groothuis-Oudshoorn, *Journal of statistical software* **45**, 1 (2011).
- F. Pedregosa, G. Varoquaux, A. Gramfort, V. Michel, B. Thirion, O. Grisel, M. Blondel, P. Prettenhofer, R. Weiss, V. Dubourg, et al., *the Journal of machine Learning research* **12**, 2825 (2011).
- A. S. Shirkhorshidi, S. Aghabozorgi, and T. Y. Wah, *PLoS ONE* **10**, e0144059 (2015).
- S. Wang, J. G. Yabes, and C.-C. H. Chang, *Journal of Data Science* **19**, 15 (2021).
- A. Shifaz, C. Pelletier, F. Petitjean, and G. I. Webb, *Knowledge and Information Systems* **65**, 2665 (2023).
- F. R. Chung and L. Lu, *Complex graphs and networks*, 107 (American Mathematical Soc., 2006).
- J. Kim and T. Wilhelm, *Physica A: Statistical Mechanics and its Applications* **387**, 2637 (2008).
- J. B. Kruskal, *Proceedings of the American Mathematical society* **7**, 48 (1956).
- C. Song, L. K. Gallos, S. Havlin, and H. A. Makse, *Journal of Statistical Mechanics: Theory and Experiment* **2007**, P03006 (2007).
- C. Schölzel, “Nonlinear measures for dynamical systems (nolds),” <https://github.com/CSchoel/nolds> (2019).
- Y. M. Omar and P. Plapper, *Entropy* **22**, 1417 (2020).
- J. King and N. Papitashvili, *Journal of Geophysical Research: Space Physics* **110** (2005), 10.1029/2004JA010649.
- I. Richardson and H. Cane, *Solar Physics* **270**, 609 (2011).
- A. Mishev, N. Larsen, E. Asvestari, A. Sáiz, M. A. Shea, D. T. Strauss, D. Ruffolo, C. Banglieng, S. Seunarine, M. L. Duldig, et al., *Advances in Space Research* **74**, 4160 (2024).
- G. E. Box and D. R. Cox, *Journal of the Royal Statistical Society Series B: Statistical Methodology* **26**, 211 (1964).
- M. Tumminello, C. Coronnello, F. Lillo, S. Micciche, and R. N. Mantegna, *International Journal of Bifurcation and Chaos* **17**, 2319 (2007).
- Australian Bureau of Meteorology Space Weather Services, “Geophysical alerts and warnings,” <https://www.sws.bom.gov.au/Geophysical/1/4> (2025), accessed: 2025

This paper was built using the Open Journal of Astrophysics L^AT_EX template. The OJA is a journal which

provides fast and easy peer review for new papers in the **astro-ph** section of the arXiv, making the reviewing process simpler for authors and referees alike. Learn more at <http://astro.theoj.org>.