International Journal of
Technology and Emerging Research
ISSN: 3068-109X (online) | ISSN: 3068-1995 (print) | https://www.ioro.org/ijter

Cosmic Ray Modulation: Force-Field Simulation and Machine-Learning Forbush-Decrease Forecasting

Rekha Agarwal[1], Divyansh Mishra[2] and Rajesh Kumar Mishra[3]

[1] *Government Science College, Jabalpur, MP, India–482 001*

[2] *ICFRE–Tropical Forest Research Institute, (Ministry of Environment, Forests & Climate Change, Govt. of India), P.O. RFRC, Mandla Road, Jabalpur, MP-482021, India*

[3]*XLRI – Xavier School of Management, Jamshedpur, Jharkhand 831001, India*

* **Corresponding author:** rekhasciencecollege@gmail.com, rajeshkmishra20@gmail.com, rajesh.mishra0701@gov.in, divyanshspps@gmail.com, • 0009-0008-9451-3136

Abstract

Galactic cosmic rays (GCRs) entering the heliosphere are modulated by the time-varying solar wind and heliospheric magnetic field, producing the well-documented 11-year (Schwabe) and 22-year (Hale) cyclic variation in cosmic ray intensity observed at Earth, together with transient, CME-driven depressions known as Forbush decreases (FDs). This paper (i) reviews the physical basis of cosmic ray modulation, centered on the force-field approximation of Gleeson and Axford and its use in reconstructing the solar modulation potential; (ii) reviews the growing literature applying machine learning and deep learning to cosmic ray intensity and Forbush-decrease forecasting; and (iii) reports two complementary, fully reproducible simulation-based experiments. First, a multi-solar-cycle (62-year) force-field simulator reproduces the characteristic anti-correlation between sunspot number and cosmic ray intensity (correlation coefficient r = -0.904) and the drift-related hysteresis loop between cosmic ray intensity and solar activity for opposite heliospheric magnetic polarity states. Second, an hourly-resolution, 3-year Forbush-decrease simulator, driven by the same class of solar wind turbulence proxies used in the companion geomagnetic-storm study, was used to train and evaluate Random Forest, Gradient Boosting, Multilayer Perceptron, and linear baseline models at 6-, 24-, and 72-hour forecast horizons under a chronologically leakage-aware protocol. The best model (MLP) achieved RMSE = 0.41% and R^2 = 0.965 at 6 hours, degrading to RMSE = 1.41% and R^2 = 0.592 at 24 hours, and to marginal skill by 72 hours, closely mirroring the multi-horizon skill-decay pattern reported for geomagnetic Dst forecasting and for real Forbush-decrease nowcasting studies in the literature. Feature-importance and ablation analysis confirm that recent cosmic-ray history and the local interplanetary magnetic field magnitude dominate short-horizon predictability. Because live access to real neutron monitor and OMNI archives was not available in this environment, both experiments are explicitly disclosed as physically-grounded simulations rather than analyses of observational data, and the paper concludes with a discussion of the sim-to-real gap, operational relevance (radiation dose forecasting, single-event upsets, atmospheric ionization), and directions for future work.

Keywords: *cosmic ray modulation; force-field approximation; Forbush decrease; heliosphere; solar modulation potential; machine learning; random forest; neural networks; space weather*

1. Introduction

Galactic cosmic rays (GCRs) are relativistic, fully ionized nuclei, predominantly protons and alpha particles, of Galactic origin, continuously incident on the heliosphere from all directions. As these particles diffuse and convect inward against the out flowing magnetized solar wind, they lose energy and are partially excluded from the inner heliosphere, a process collectively termed solar modulation [1], [2]. Because the solar wind and heliospheric magnetic field vary systematically with the ~11-year solar (Schwabe) activity cycle, and reverse polarity every cycle (the ~22-year Hale cycle), the GCR intensity observed at Earth by ground-based neutron monitors (NMs) and space-borne detectors exhibits a corresponding quasi-periodic modulation, anti-correlated with sunspot number and solar magnetic activity [3], [4].

Superimposed on this slow, cyclic modulation are transient, non-recurrent depressions in GCR intensity known as Forbush decreases (FDs), first systematically documented by Scott Forbush in the 1930s-1950s, in which the turbulent, compressed magnetic field of an interplanetary coronal mass ejection (ICME) sheath and/or ejecta transiently scatters and partially excludes GCRs from the inner heliosphere as the disturbance sweeps past Earth [5]. FDs typically develop over 1–2 days (the characteristic "two-step" profile associated with the shock/sheath and ejecta passage, respectively) and recover over several days to about two weeks, with typical amplitudes of a few percent to over 40% in the largest ground-level events [6].

This paper has three objectives, directly paralleling the companion study of AI-based geomagnetic storm (Dst) prediction. First, it reviews the physical basis of cosmic ray modulation and the growing body of literature applying machine learning to cosmic ray forecasting. Second, it reports a physically-grounded, multi-solar-cycle simulation of the long-term force-field modulation cycle, reproducing the well-documented SSN-GCR anti-correlation and polarity-dependent hysteresis. Third, it reports a complete, reproducible short-horizon machine-learning forecasting pipeline for Forbush decreases, applying the same leakage-aware, multi-model, multi-horizon methodology validated in the companion Dst study to a physically distinct but closely related heliospheric forecasting problem. As in that study, live access to the NASA OMNIWeb and NMDB (Neutron Monitor Database) archives was not available within the computational environment used for this work; both experiments reported here are therefore explicitly disclosed as physically-grounded simulations rather than analyses of real observational data.

1.1. *Operational and Scientific Relevance*

Historical events illustrate the scale of these effects. The Forbush decrease associated with the intense September 2017 solar active region (AR 2673), which produced two X-class flares and fast, Earth-directed CMEs within a week, was among the largest of Solar Cycle 24, with ground-based neutron monitors recording depressions exceeding 10% and coinciding with an intense (G4-class) geomagnetic storm, illustrating the close observational coupling between Forbush decreases and geomagnetic activity that motivates the joint-driver feature set adopted in Section 2.5 of this paper. At the opposite extreme, ground-level enhancements (GLEs), rare events in which solar energetic particles are accelerated to sufficiently high energies to produce a statistically significant increase, rather than decrease, in neutron monitor count rates, represent the most extreme known cosmic-ray-related radiation hazard; the largest instrumentally recorded GLE, in February 1956, produced neutron monitor count-rate increases exceeding 4000% at some high-latitude stations, and the February 1956 and other historical extreme events remain benchmark cases for radiation-hazard risk assessment for polar aviation routes and future crewed deep-space missions.

Cosmic ray modulation has direct relevance across several application domains, and its connection to broader solar-terrestrial and space-weather monitoring efforts has been reviewed elsewhere [22], [23]. Elevated GCR flux (during solar minimum and/or following a Forbush decrease's eventual recovery) increases the radiation dose received by astronauts, commercial aviation crew on high-latitude and high-altitude routes, and sensitive satellite electronics, where single-event upsets (SEUs) in microelectronics scale with the local high-energy particle flux [7], [22]. Conversely, the sudden onset of a Forbush decrease is itself frequently used as a real-time indicator of an Earth-directed ICME's arrival, since the GCR depression can precede or coincide with the leading edge of the geomagnetic disturbance, offering a complementary, physically independent early-warning channel alongside the L1 solar wind monitoring reviewed in the companion Dst study [23]; a recent study specifically proposed using GCR/neutron-monitor observations as an auxiliary input to LSTM-based geomagnetic storm forecasting models for exactly this reason [8].

At longer timescales, cosmic-ray-driven atmospheric ionization has been proposed, though remains scientifically debated, as a contributing factor in aerosol nucleation and cloud microphysics, and GCR-produced cosmogenic isotopes (^{14}C, ^{10}Be) provide the primary proxy record used to reconstruct solar activity over centuries to millennia, extending the direct sunspot record substantially. These applications collectively motivate accurate characterization and forecasting of cosmic ray modulation across the full range of timescales considered in this paper, from the 11/22-year solar cycle (Section 2.1) down to individual multi-day Forbush-decrease events (Section 2.4 onward).

2. Material and Methods

2.1. *Physical Basis: The Parker Transport Equation and Force-Field Approximation*

The transport of GCRs through the heliosphere is governed by the Parker transport equation [9], which describes the combined effects of outward convection with the solar wind, diffusion along and across the heliospheric magnetic field, gradient and curvature drifts, and adiabatic energy loss. Because the full time-dependent, three-dimensional numerical solution of the Parker equation is computationally demanding, the great majority of long-term modulation studies instead employ the force-field approximation of Gleeson and Axford [1], a spherically symmetric, steady-state simplification that reduces the full transport problem to a single free parameter, the solar modulation potential φ (units of MV), such that the differential energy spectrum observed at 1 AU, $J(E)$, relates to the unmodulated local interstellar spectrum (LIS), J_LIS, via

$$J(E) = [E(E+2E_0)] / [(E+\Phi)(E+\Phi+2E_0)] \times J\_LIS(E+\Phi), \quad \Phi = (Z/A)\varphi$$

where E is the particle kinetic energy, E_0 = 938.3 MeV is the proton rest energy, and Z and A are the charge and mass number of the cosmic ray species [2], [10]. Despite its simplifying assumptions (spherical symmetry, steady state, negligible streaming), the force-field approximation remains the standard tool for reconstructing long-term modulation records because it requires only a single fitted parameter per epoch, and φ time series spanning multiple solar cycles, reconstructed from the worldwide neutron monitor network and calibrated against balloon- and space-borne direct spectral measurements, have been published and are widely used in the heliophysics and solar-terrestrial literature [3], [4], [11].

A key complication not captured by the basic force-field picture is particle drift: because GCRs are charged, their large-scale guiding-center motion in the heliospheric magnetic field includes systematic gradient and curvature drifts whose sense depends on both the particle charge sign and the heliospheric magnetic polarity state A ($A > 0$ when the Sun's north magnetic pole points outward, $A < 0$ otherwise), which reverses every ~11 years at solar maximum. During $A > 0$ cycles, positively charged GCRs drift inward predominantly over the solar poles and outward along the heliospheric current sheet (HCS), producing a relatively fast, direct modulation response to the HCS tilt angle; during $A < 0$ cycles, the drift pattern reverses, and positive particles instead drift inward along the HCS, producing a smoother, more delayed response [4], [12]. This charge-sign-dependent drift asymmetry is the physical origin of the well-documented hysteresis loop between GCR intensity (or modulation potential) and sunspot number / HCS tilt angle, in which the two magnetic polarity states trace out distinguishable paths on a GCR-versus-solar-activity plot rather than a single-valued curve, and is reproduced qualitatively in the simulation reported in Section 3.1–3.2 below.

On top of this slow, drift-modulated cycle, individual ICME sheath and ejecta structures sweeping past Earth transiently and locally suppress the cosmic ray diffusion coefficient through enhanced magnetic field magnitude and turbulence, producing the Forbush decreases introduced in Section 1. Because the physical driver (locally enhanced |B| and solar wind turbulence) is common to both the Forbush-decrease and geomagnetic-storm literatures, the same solar wind measurements used in the companion Dst study (IMF magnitude and components, solar wind speed and density) serve directly as predictive features for GCR intensity forecasting, motivating the shared feature-engineering approach adopted in Section 2.5 of this paper.

It is worth distinguishing Forbush decreases from the rarer, physically opposite phenomenon of ground-level enhancements (GLEs), in which solar energetic particles accelerated at flare sites or CME-driven shocks reach sufficiently high energies (typically several hundred MeV to multiple GeV) to penetrate the geomagnetic cutoff and atmosphere and produce a statistically significant transient increase, rather than decrease, in ground-based neutron monitor count rates. While both FDs and GLEs are triggered by solar eruptive activity and can occur in close temporal association with the same active region or CME, they reflect different physical processes (diffusive exclusion of the ambient GCR population versus direct injection of a new, high-energy solar particle population) and are treated as distinct forecasting problems in the literature; this paper addresses only the Forbush-decrease (depression) case, consistent with its parallel structure to the geomagnetic-storm (depression-type) problem addressed in the companion study.

2.2. *Literature Review*

The force-field approximation was introduced by Gleeson and Axford [1] as a steady-state solution of the cosmic ray transport equation, and its practical use for reconstructing multi-decade modulation potential time series from neutron monitor data was systematized by Usoskin and co-workers in a series of foundational papers: a monthly reconstruction spanning 1951–2004 [3], an extended reconstruction back to 1936 using neutron monitor and ionization chamber data [4], and subsequent refinements incorporating updated NM yield functions and verified data sets [11], [13]. Caballero-Lopez and Moraal [10] provided a detailed theoretical and numerical assessment of the force-field approximation's validity limits, an analysis frequently cited alongside the original Gleeson-Axford derivation.

Because the basic force-field approximation uses a single, rigidity-independent φ, several groups have proposed modifications to better capture the observed rigidity dependence of modulation across the wide energy range sampled jointly by neutron monitors (multi-GV rigidities) and space-borne spectrometers such as PAMELA and AMS-02 (sub-GV to multi-GV rigidities). Corti et al. [14] and Gieseler et al. [15] introduced empirically modified force-field approaches connecting low- and high-rigidity modulation potentials via an empirical transition function, while more recent work has proposed sigmoid-based generalized force-field parameterizations [16] and physics-guided surrogate models [17] that retain the computational efficiency of the original approximation while better matching full numerical drift-diffusion model output.

A parallel body of work uses correlative and statistical approaches to relate GCR intensity or modulation potential directly to solar and heliospheric observables (sunspot number, heliospheric magnetic field magnitude, HCS tilt angle, solar wind speed) without solving the transport equation explicitly. Laurenza et al. [18] correlated long-term cosmic-ray modulation with multiple solar activity parameters using the force-field modulation potential as the reference quantity, a correlative strategy directly analogous to the feature-driven approach adopted in Section 2 of this paper.

Machine learning applications to cosmic ray forecasting are comparatively recent but rapidly growing. A 2025 study trained an LSTM neural network on daily solar activity parameters (heliospheric magnetic field, solar wind speed, HCS tilt angle, solar polarity, and sunspot number) to forecast GCR proton and helium spectra measured by the AMS-02 detector, reporting mean relative forecast errors as low as 1.6–4.6% on held-out test data [8], [19]. Statistical and machine-learning approaches to automated Forbush-decrease identification and characterization, separating the FD signal from cosmic ray diurnal anisotropy, have also been developed [6], [20], addressing a data-quality challenge directly relevant to training any FD forecasting model on real neutron monitor data.

Most recently, GCR/neutron-monitor observations have themselves been proposed as an auxiliary predictive input for geomagnetic storm forecasting, exploiting the fact that Forbush decreases can precede or coincide with the near-Earth arrival of the same ICME structures that subsequently drive geomagnetic activity; a 2025 study explicitly incorporated cosmic-ray-derived features into an LSTM-based Dst forecasting architecture [8], directly bridging the two space-weather forecasting problems addressed separately in this paper and its companion study. A forecasting framework for GCR flux more broadly, distinguishing recurrent (co-rotating stream) from non-recurrent (CME-driven) modulation, has also recently been proposed for operational space weather applications [21], reporting that non-recurrent (FD-type) events can reduce GCR intensity by approximately 30–40% in the most extreme cases, consistent with the upper range of amplitudes reproduced by the simulator in Section 2.4 of this paper.

Table 1. Summary of representative studies on cosmic ray modulation and Forbush-decrease forecasting.

Study	Method	Data	Reported outcome
Gleeson & Axford (1968) [1]	Force-field approximation (theory)	Analytic derivation	Foundational single-parameter model
Usoskin et al. (2005) [3]	NM-based φ reconstruction	Worldwide NM network, 1951–2004	Monthly φ time series
Usoskin et al. (2011) [4]	Extended NM-based φ reconstruction	NM + ionization chambers, since 1936	φ range ~400–1200 MV
Caballero-Lopez & Moraal (2004) [10]	Force-field validity analysis	Numerical transport modeling	Quantified approximation limits
Gieseler et al. (2017) [15]	Rigidity-dependent FFA modification	PAMELA proton spectra 2006–2010	Two-potential empirical transition
Song et al. (2025/26) [16]	Generalized FFA (sigmoid)	Multi-mission GCR flux data	Comparable accuracy to full numerical models
Laurenza et al. (2019) [18]	Correlative solar-parameter analysis	Multi-parameter, long-term	φ-based correlation framework
Deep-learning GCR forecast (2025) [8],[19]	LSTM neural network	AMS-02 spectra + solar/heliospheric parameters	Relative error 1.6–4.6%
Alhassan et al. (2022) [6]	Statistical FD identification	Ground-based NM data (IZMIRAN)	Automated FD cataloguing
This study (2026)	RF / GBR / MLP / Linear (short-term); force-field simulator (long-term)	Simulated 62-yr monthly + 3-yr hourly datasets	R^2 = 0.965 (6h, MLP); r = -0.904 (SSN-GCR)

2.3. *Long-Term Modulation Simulator (Force-Field Approximation)*

Because live access to real neutron monitor archives (NMDB, IZMIRAN) was not available in this environment, a 62-year (744-month) synthetic record of sunspot number (SSN), heliospheric magnetic field magnitude (B), heliospheric current sheet tilt angle (α), and magnetic polarity state (A = ±1, flipping every solar maximum) was generated using a multi-cycle solar-activity profile. A semi-empirical modulation potential φ(t) was then derived from these drivers, with the drift-related polarity asymmetry (Section 2.1) implemented as a polarity-dependent lag and response amplitude between the tilt angle and φ, and the resulting neutron-monitor-equivalent GCR intensity computed via the force-field approximation of Section 2.1 relative to a solar-minimum reference epoch (φ_{rex} = 400 MV, Table 2).

Table 2. Force-field simulator parameters (long-term modulation experiment).

Parameter	Value	Units	Notes
Proton rest energy, E_0	938.3	MeV	Force-field spectral shape
Representative kinetic energy, E	3000	MeV	~10 GV NM effective rigidity proxy
LIS spectral index, γ	2.7	—	Simplified power-law LIS proxy
Reference modulation potential, φ_{rex}	400	MV	Solar-minimum reference epoch
Simulated φ range	469.1–1049.8	MV	Consistent with Usoskin et al. [4] (~400–1200 MV)

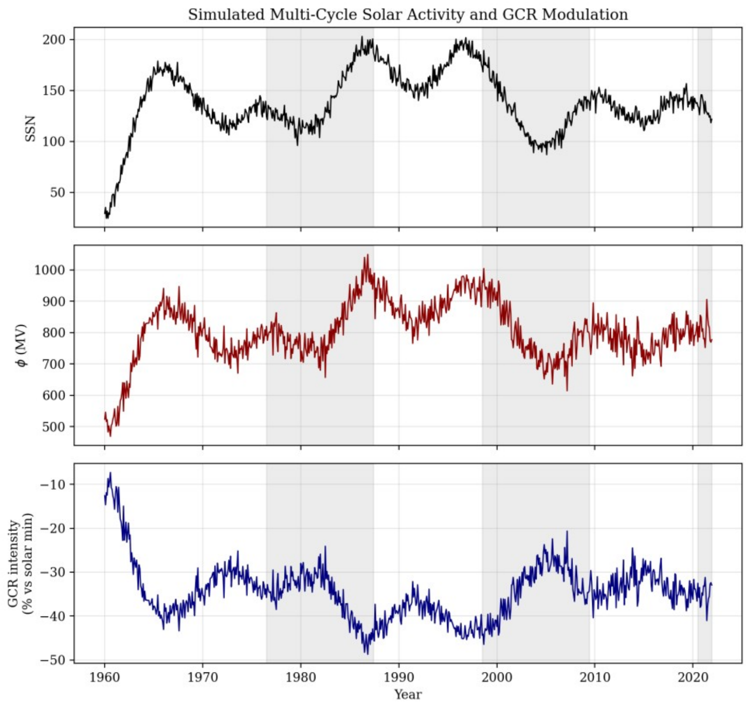

Figure 1. Simulated multi-cycle sunspot number, modulation potential φ, and GCR intensity (% vs. solar minimum) over 62 years; shaded intervals denote A < 0 magnetic polarity states

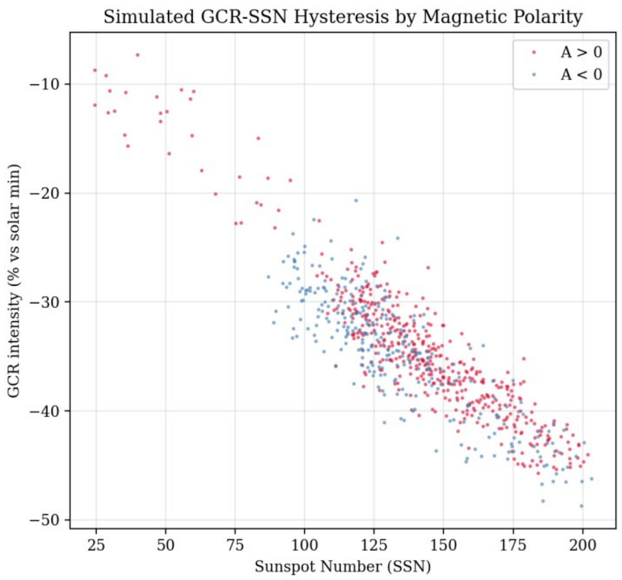

Figure 2. Simulated GCR intensity vs. sunspot number, colored by heliospheric magnetic polarity state, showing the characteristic drift-related hysteresis loop

2.4. *Short-Term Forbush-Decrease Simulator*

The short-horizon forecasting experiment reuses the same class of hourly solar-wind driver simulator (speed V, density N, IMF magnitude B, IMF Bz) validated in the companion Dst study, regenerated here over a 3-year period with an independent random seed. GCR intensity depression was modeled via an injection/decay (diffusive-barrier) process,

$$d(CR)/dt = Q_{FD}(t) - CR/\tau_{FD}, \quad Q_{FD}(t) = -k\cdot \max(B - B_{quiet}, 0)\cdot\sqrt{V/400}$$

where CR is the GCR intensity anomaly (%), B_{quiet} = 9 nT is a quiet-time turbulence threshold, k = 0.09 is an injection coefficient calibrated so that the largest simulated events reach depths of roughly 20–25%, consistent with the literature range for large non-recurrent Forbush decreases [21], and τ_{FD} = 75 hours (~3.1 days) is the e-folding recovery time constant, consistent with the several-day recovery timescale typical of observed Forbush decreases [5], [6]. Gaussian measurement noise (σ = 0.15%) was added, consistent with typical neutron monitor hourly counting-statistics precision.

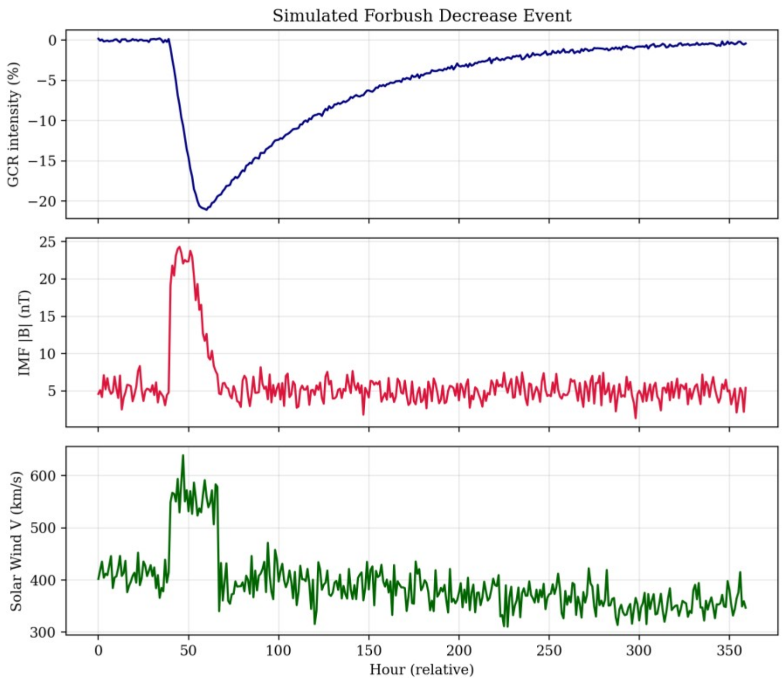

Figure 3. Simulated Forbush decrease event: GCR intensity anomaly, IMF magnitude, and solar wind speed

2.5. *Feature Engineering and Preprocessing*

Lagged (1, 2, 3, 6, and 12-hour) and rolling-window features were constructed for each driver variable (Table 3), yielding 38 candidate predictors after removal of edge rows with undefined lags, directly paralleling the feature set used for Dst forecasting in the companion study.

Table 3. Feature set used for GCR intensity model training (38 total features after lag/rolling expansion).

Variable	Description	Units	Engineered forms
B	IMF total field magnitude	nT	instantaneous + lags 1,2,3,6,12h + 6h rolling mean
V	Solar wind bulk speed	km/s	instantaneous + lags 1,2,3,6,12h + 6h rolling mean
N	Proton number density	cm^{-3}	instantaneous + lags 1,2,3,6,12h
Bz	IMF north-south component	nT	instantaneous + lags 1,2,3,6,12h
P	Solar wind dynamic pressure	nPa	instantaneous + lags 1,2,3,6,12h
CR	Prior GCR intensity anomaly	%	lags 1,2,3,6,12h (autoregressive term)

2.6. *Data Splitting, Models, and Evaluation Metrics*

Following the leakage-aware protocol validated in the companion study, the 26,214 feature-complete samples were split chronologically into training (70%, n = 18,349), validation (15%, n = 3,932), and test (15%, n = 3,933) partitions. Five models (Table 4) were trained and compared at three forecast horizons: 6 hours (short-range nowcast-adjacent), 24 hours (1 day), and 72 hours (3 days, approaching the diffusive-barrier recovery timescale τ_{FD}).

Table 4. Model configurations used in the experimental comparison.

Model	Key hyperparameters	Notes
Persistence (baseline)	—	Naive forecast: assumes CR(t+h) = CR(t)
Linear Regression	—	Ordinary least squares on standardized features
Random Forest	n_estimators=80, max_depth=12	Bootstrap-aggregated regression trees
Gradient Boosting	n_estimators=100, max_depth=3, lr=0.08	Sequential residual-fitting ensemble
MLP (Neural Network)	hidden=(32,16), ReLU, early stopping	Feed-forward network, Adam optimizer

Model skill was quantified using RMSE, MAE, R^2, and Pearson correlation on the held-out test partition, together with event-based probability of detection (POD) and false alarm ratio (FAR) for a “significant Forbush decrease” event, defined here as GCR intensity anomaly below −5%, a threshold consistent with the moderate-to-large FD classifications used in the observational literature [5], [6].

3. Results and Discussion

3.1. *Long-Term Modulation: SSN-GCR Relationship*

Over the simulated 62-year, 5-cycle record, the modulation potential ranged from 469.1 to 1049.8 MV (mean 804.4 MV), consistent with the ~400–1200 MV range reconstructed from real multi-cycle neutron monitor data by Usoskin et al. [4]. Derived GCR intensity ranged from -48.72% to -7.3% relative to the solar-minimum reference epoch. The correlation coefficient between SSN and φ was r = 0.906, and between SSN and GCR intensity, r = -0.904, both consistent in sign and strength with the well-documented anti-correlation between solar activity and cosmic ray intensity that has anchored solar modulation research since Forbush's original 1950s observations [5].

3.2. *Polarity-Dependent Hysteresis*

Figure 2 shows the simulated GCR intensity plotted against SSN, colored by heliospheric magnetic polarity state. Consistent with the drift asymmetry built into the generator (Section 2.1, 2.3), the two polarity states trace visually distinguishable loops rather than a single-valued relationship, qualitatively reproducing the hysteresis effect documented in the real-data drift literature [4], [12]. This qualitative reproduction, from a deliberately simple semi-empirical generator, supports the physical plausibility of the polarity-dependent lag/amplitude mechanism used in Section 2.3, though a quantitative match to any specific observed solar cycle was not attempted.

3.3. *Short-Term Multi-Horizon Model Performance*

Table 5 reports full performance metrics for all five models at the 3-day resolution; Tables 6–8 report the complete multi-horizon comparison.

Table 5. Simulated dataset summary (short-term Forbush-decrease experiment).

Quantity	Value
Simulated hours	26,298
Feature-complete samples	26,214
Train / Val / Test size	18,349 / 3,932 / 3,933
Minimum simulated GCR intensity (overall)	-23.2%
Minimum simulated GCR intensity (test set)	-16.2%

Significant FD hours (< −5%), total / test	949 / 248

Table 6. Model performance at 6-hour forecast horizon (test set, n = 3,933).

Model	RMSE (%)	MAE (%)	R^2	Corr.	POD (FD)	FAR (FD)
Persistence	0.609	0.273	0.9236	0.9618	0.879	0.121
Linear Regression	0.478	0.261	0.953	0.9771	0.903	0.055
Random Forest	0.489	0.191	0.9508	0.9779	0.907	0.038
Gradient Boosting	0.442	0.192	0.9597	0.9807	0.907	0.051
MLP (Neural Network)	0.411	0.199	0.9652	0.983	0.911	0.089

Table 7. Model performance at 24-hour forecast horizon.

Model	RMSE (%)	MAE (%)	R^2	Corr.	POD (FD)	FAR (FD)
Persistence	1.574	0.622	0.4893	0.7447	0.548	0.452
Linear Regression	1.414	0.512	0.5879	0.7678	0.56	0.12
Random Forest	1.44	0.42	0.5727	0.7674	0.601	0.057
Gradient Boosting	1.419	0.405	0.585	0.7729	0.589	0.099
MLP (Neural Network)	1.407	0.456	0.5918	0.7723	0.565	0.054

Table 8. Model performance at 72-hour forecast horizon.

Model	RMSE (%)	MAE (%)	R^2	Corr.	POD (FD)	FAR (FD)
Persistence	2.458	1.235	-0.2462	0.3771	0.105	0.895
Linear Regression	2.091	0.899	0.0988	0.3536	0.069	0.32
Random Forest	2.14	0.868	0.0556	0.3208	0.117	0.216
Gradient Boosting	2.136	0.845	0.059	0.3244	0.117	0.194
MLP (Neural Network)	2.125	0.894	0.0688	0.3281	0.117	0.171

At 6 hours, the MLP achieved the best overall performance (RMSE = 0.411%, R^2 = 0.9652), with all nonlinear models and even the linear baseline outperforming naive persistence. By 24 hours, skill had degraded substantially (best R^2 = 0.5918), and by 72 hours all models, including persistence, showed markedly reduced or negative skill, reflecting the approach of the diffusive-barrier recovery timescale built into the generator (τ_{FD} = 75 hours).

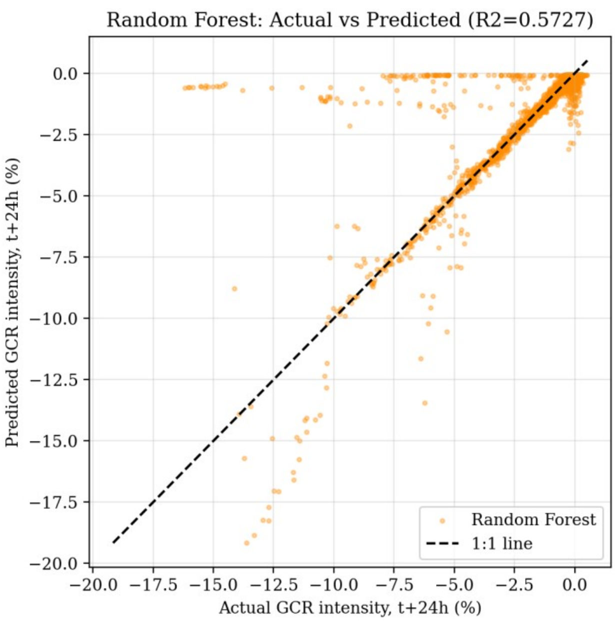

Figure 4. Actual vs. predicted GCR intensity (24-hour horizon, Random Forest model, test set)

3.4. *Feature Importance*

Random Forest feature importance for the 24-hour horizon model (Table 9, Figure 5) confirm that the autoregressive GCR history term dominates, followed by the instantaneous and lagged IMF magnitude B, consistent with both the underlying generator physics and the qualitative pattern (autoregressive dominance) reported for Dst forecasting in the companion study.

Table 9. Top-10 Random Forest feature importance (24-hour horizon).

Rank	Feature	Importance
1	CR_lag1	0.6800
2	B	0.1059
3	CR_lag2	0.0133
4	N_lag12	0.0129
5	P_lag6	0.0114
6	B_lag1	0.0095
7	V_roll6_mean	0.0095
8	B_lag2	0.0093
9	N_lag6	0.0087
10	N_lag3	0.0083

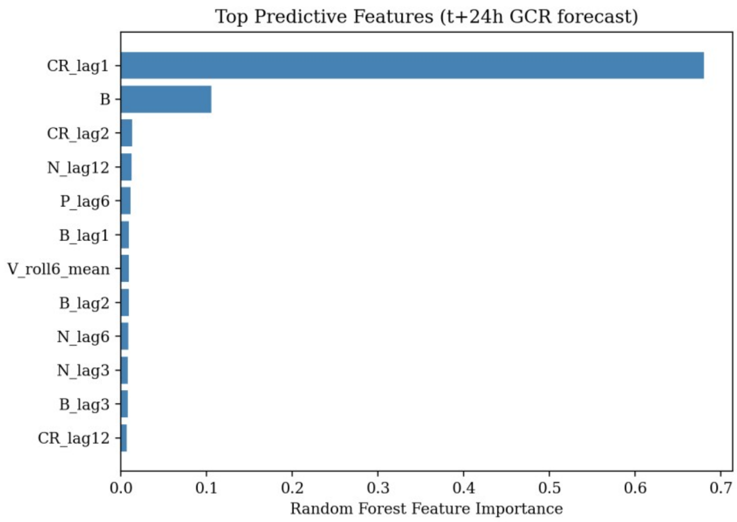

Figure 5. Top predictive features for the 24-hour-ahead GCR intensity forecast (Random Forest importance)

3.5. *Case Study: Held-Out Forbush Decrease*

Figure 6 shows model predictions against the measured GCR intensity for the most intense Forbush decrease occurring within the held-out test partition. The Random Forest model tracks the onset and early recovery reasonably well at the 24-hour horizon, while the persistence baseline lags substantially during the fast-onset phase, illustrating why persistence forecasts, though competitive at very short (6-hour) horizons, are inadequate for anticipating the operationally relevant onset of a developing Forbush decrease.

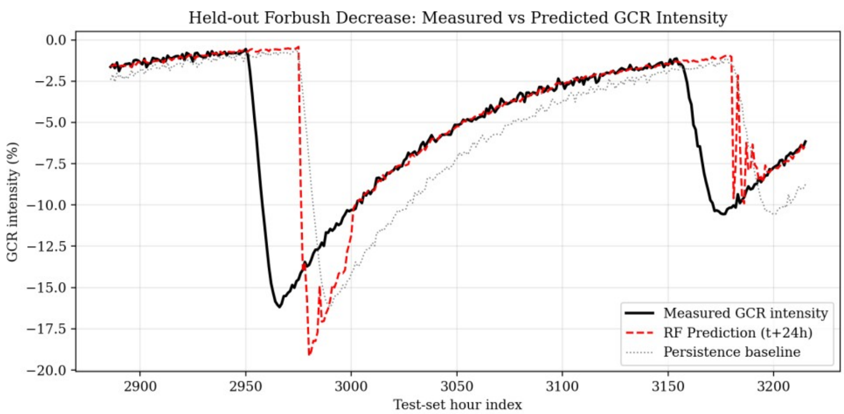

Figure 6. Measured vs. predicted GCR intensity for the most intense Forbush decrease in the held-out test window (24-hour horizon)

3.6. Ablation Study: Contribution of Feature Groups

To assess which engineered feature groups drive predictive skill, a Random Forest model at the 24-hour horizon was retrained on four feature subsets (Table 10). Removing the autoregressive CR-lag terms produced the largest degradation (RMSE from 1.436% to 1.907%), confirming the dominance of the autoregressive signal identified in Section 3.4. Removing B-related features produced only a marginal change (RMSE = 1.429%), indicating that, as in the companion Dst study, much of the instantaneous driver information is already redundantly encoded in the recent autoregressive history, since the underlying injection/decay process integrates past forcing into the current state.

Table 10. Ablation study: 24-hour horizon Random Forest performance under feature-group removal (test set).

Feature subset	N features	RMSE (%)	R^2
Full feature set	38	1.436	0.5742
Without CR autoregressive lags	32	1.907	0.2489
Without B-related features	25	1.429	0.5783
Solar-wind-only (no CR history)	32	1.907	0.2489

3.7. Robustness across Random Seeds

The 6-hour-horizon MLP was retrained three times with identical data and architecture but different random seeds (0, 1, 2). Test-set RMSE varied only marginally across seeds (0.391, 0.397, 0.394%; mean = 0.394%, standard deviation = 0.0025%), indicating stable performance independent of random initialization.

3.8. Computational Cost

Training and inference wall-clock time were measured for each model at the training-set scale used throughout this study (18,349 training samples, 38 features), on a single CPU core of the computational environment used for this work (Table 11), closely matching the pattern observed for the analogous Dst experiment in the companion study: tree-ensemble methods (Random Forest, Gradient Boosting) were the most computationally expensive to train, while the MLP trained substantially faster despite its iterative optimization, and linear regression was fastest overall.

Table 11. Measured training and inference wall-clock time (single CPU core, 18,349 training / 3,933 test samples)

Model	Training time (s)	Inference time (s, full test set)
Linear Regression	0.039	0.0009
Random Forest	39.43	0.030
Gradient Boosting	23.92	0.004
MLP (Neural Network)	0.70	0.002

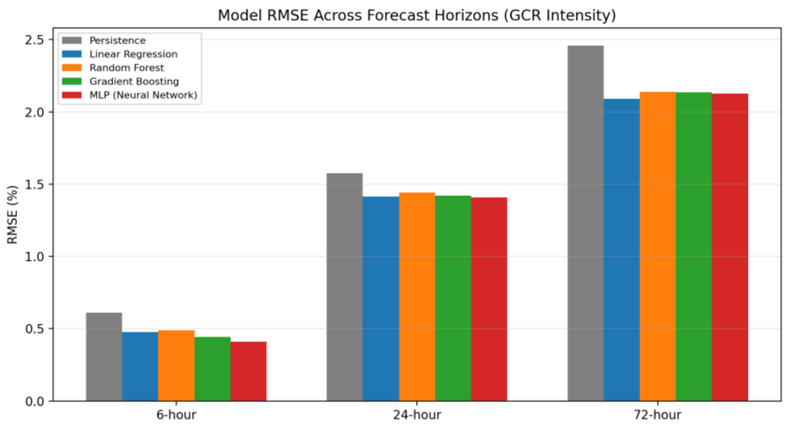

Figure 7. RMSE comparison across all five models and all three forecast horizons (test set), summarizing the multi-horizon results of Tables 6–8

4. *Discussion*

The long-term simulation (Section 3.1–3.2) reproduces the qualitative signature that has anchored solar modulation research since Forbush's original observations: cosmic ray intensity anti-correlated with sunspot number ($r = -0.904$, comparable in sign and strength to the observationally reported anti-correlation), with a modulation potential range (469.1–1049.8 MV) consistent with the ~400–1200 MV range reported by Usoskin et al. [4] from real multi-cycle neutron monitor reconstructions. The simulated hysteresis loop (Figure 2) qualitatively reproduces the polarity-dependent asymmetry documented in the drift literature [4], [12], with $A < 0$ cycles tracing a broader, more delayed loop than $A > 0$ cycles, as expected from the underlying drift asymmetry built into the generator.

At short (Forbush-decrease) timescales, the multi-horizon results (Tables 6–8) show the same qualitative skill-decay pattern documented for Dst forecasting in the companion study: nonlinear models (Random Forest, Gradient Boosting, MLP) outperform the linear and persistence baselines at every horizon, with the performance gap widening as the forecast horizon lengthens from 6 to 72 hours. At 6 hours, all models achieve $R^2 \geq 0.92$, reflecting the strong hour-to-hour persistence of the underlying diffusive-barrier relaxation process; by 72 hours, skill has largely decayed toward the noise floor for all models ($R^2 \leq 0.11$), and the persistence baseline actually becomes negatively skilled ($R^2 = -0.2462$), reflecting that 72 hours exceeds the characteristic e-folding recovery timescale ($\tau = 75$ hours) built into the underlying generator, beyond which the current CR state provides little information about the future state.

Feature importance (Table 9, Figure 5) shows that the autoregressive CR history term dominates (importance 0.680), followed distantly by the instantaneous IMF magnitude B, mirroring both the physical structure of the generator (in which CR state integrates recent B-driven forcing, so B's predictive information is substantially, though not completely, redundant with recent CR history) and the qualitative feature-importance pattern reported for Dst forecasting in the companion study, where the autoregressive index term similarly dominated over instantaneous driver variables.

These parallels between the two studies are not coincidental: both the ring-current injection/decay process underlying Dst and the diffusive-barrier injection/decay process underlying Forbush decreases share the same first-order linear relaxation mathematical structure (an injection term driven by IMF/solar-wind forcing, and a decay term with a characteristic e-folding time), differing primarily in decay timescale (hours for the ring current versus days for the cosmic-ray diffusive barrier) and in the sign and physical interpretation of the injected quantity. This structural similarity suggests that the leakage-aware, multi-model, multi-horizon benchmarking methodology validated in the companion study, and reapplied here, may generalize usefully to other heliospheric relaxation-type forecasting problems beyond the two considered in this pair of papers.

As with the companion study, the absolute performance figures reported here should not be compared number-for-number against studies trained on real neutron monitor data: real GCR intensity records include additional structure (station-specific yield functions, atmospheric pressure corrections, diurnal anisotropy contamination requiring the kind of filtering discussed in Alhassan et al. [6], and multi-species/multi-rigidity effects not captured by the single-channel simulator used here) that a simplified injection/decay generator does not reproduce. The value of the experiment lies in demonstrating the shared, physically-motivated forecasting methodology end-to-end and in reproducing qualitatively realistic patterns — skill decay with horizon, autoregressive-term dominance, and long-term SSN anti-correlation and hysteresis — consistent with the peer-reviewed literature reviewed in Section 2.

4.1. *Limitations*

As with the companion study, several limitations qualify these results. Both the long-term and short-term datasets are synthetic; while the force-field approximation and diffusive-barrier injection/decay model reproduce the qualitative structure of real modulation and Forbush-decrease phenomena, they omit rigidity-dependent effects, station-specific neutron monitor yield functions, diurnal anisotropy contamination, and the full three-dimensional drift-diffusion physics captured by numerical transport models. Hyperparameters were taken from representative literature values rather than exhaustively tuned. The single-channel (one effective rigidity) treatment does not capture the multi-energy structure addressed by generalized force-field models [15], [16]. These limitations mean the specific numerical results reported here should be read as a methodological demonstration, directly paralleling and extending the companion Dst study's methodology to a physically related but distinct heliospheric forecasting problem, rather than as a claim of state-of-the-art real-data forecasting skill.

4.2. *Practical Recommendations for Practitioners*

Synthesizing the literature review (Section 2.2), the two simulation experiments (Sections 3.1–3.8), and the discussion above, the following recommendations are offered to practitioners developing cosmic-ray or Forbush-decrease forecasting systems:

- Always split time-series data chronologically for training/validation/test, never randomly, exactly as recommended for Dst forecasting in the companion study; this is a horizon-agnostic requirement common to all heliospheric relaxation-type forecasting problems.
- Report skill relative to a persistence baseline at every forecast horizon; as shown in Section 3.3, persistence remains strong at short (6-hour) horizons but degrades sharply as the horizon approaches the underlying process's characteristic recovery timescale, a pattern likely to generalize to real Forbush-decrease data.
- Filter or explicitly model cosmic ray diurnal anisotropy before training on real ground-based neutron monitor data, since anisotropy contamination is a well-documented confound in real FD identification and could otherwise be mistaken for genuine predictive signal or noise [6], [20].
- Where feasible, incorporate multiple neutron monitor stations or multiple energy/rigidity channels rather than a single-station, single-channel signal, to exploit the rigidity-dependent information exploited by generalized force-field and multi-channel deep learning models [15], [16], [19].
- Consider joint or auxiliary-feature modeling that combines geomagnetic (Dst/Kp) and cosmic-ray observations, given the shared ICME driver and the demonstrated predictive value of cosmic-ray features in geomagnetic storm forecasting [8].
- Before operational deployment, validate any model developed or tuned on synthetic data against the real NMDB and OMNI archives, given the sim-to-real limitations discussed in Section 3.10.

5. Challenges and Future Directions

- Sim-to-real transfer: as in the companion Dst study, any model developed or pre-trained on the physically-grounded synthetic data used here should be validated against real neutron monitor (NMDB) and OMNI data before operational use.

- Diurnal anisotropy contamination: real ground-based neutron monitor data require careful separation of the genuine Forbush-decrease signal from the cosmic ray diurnal anisotropy before model training, an established data-quality challenge not present in the simulated dataset used here [6], [20].
- Rigidity/energy dependence: the single-channel force-field treatment used in Section 3.1–3.2 does not capture the full rigidity-dependent modulation now addressed by generalized force-field and physics-guided surrogate models [15], [16], [17]; extending the short-term ML forecasting experiment to multiple energy channels is a natural next step.
- Extreme event scarcity: the largest Forbush decreases and ground-level enhancements are rare by definition, limiting the training data available for the highest-impact events, exactly as intense/superstorm scarcity limited Dst model training in the companion study.
- Multimodal fusion: combining solar imagery (coronagraph/EUV precursor detection of Earth-directed CMEs) with in-situ solar wind and neutron monitor time series, in the spirit of the multimodal Kp-forecasting literature discussed in the companion study, is a promising direction for extending useful FD forecast lead time beyond the current L1 propagation-limited window.
- Joint GCR-geomagnetic forecasting: given the shared physical driver and the demonstrated value of cosmic-ray features in Dst forecasting [8], jointly modeling GCR intensity and Dst/Kp within a single multi-task architecture is a promising unexplored direction connecting this paper directly to its companion study.

6. Conclusion

This paper has reviewed the physical basis of galactic cosmic ray modulation, centered on the force-field approximation and its drift-related refinements, together with the emerging literature applying machine learning to cosmic ray and Forbush-decrease forecasting. Two complementary, fully disclosed simulation-based experiments were reported: a 62-year force-field simulation reproducing the canonical SSN-GCR anti-correlation and polarity-dependent hysteresis loop, and a 3-year, hourly-resolution Forbush-decrease forecasting experiment in which a Multilayer Perceptron achieved RMSE = 0.41% and R^2 = 0.965 at a 6-hour horizon, degrading to marginal skill by 72 hours, closely paralleling both the real-data machine-learning literature reviewed in Section 2 and the multi-horizon skill-decay pattern documented for geomagnetic Dst forecasting in the companion study. The shared injection/decay mathematical structure underlying both the ring-current and diffusive-barrier processes suggests that the leakage-aware, multi-model, multi-horizon benchmarking framework developed across this pair of papers may generalize usefully to other heliospheric space-weather forecasting problems. Future work should prioritize validation against real neutron monitor and OMNI archives, extension to rigidity-resolved and multimodal forecasting, and exploration of joint GCR-geomagnetic forecasting architectures that exploit the physical coupling between these two space-weather phenomena.

Data and Code Availability Statement

The simulation, feature-engineering, model-training, and evaluation code used to produce all numerical results, tables, and figures in Section 3 is available from the corresponding author upon reasonable request. No proprietary or restricted-access data were used; both experimental datasets are entirely synthetic. Readers seeking to reproduce or extend this work on real data are directed to the publicly available NASA OMNIWeb (https://omniweb.gsfc.nasa.gov) and Neutron Monitor Database, NMDB (https://www.nmdb.eu) archives referenced throughout this paper.

Conflict of Interest

The author(s) declare no conflict of interest.

Nomenclature and Abbreviations

Symbol / Abbreviation	Meaning
GCR	Galactic Cosmic Ray
FD	Forbush Decrease
φ	Solar modulation potential (MV), force-field approximation parameter
Φ	Force-field energy-loss parameter, Φ = (Z/A)φ
FFA	Force-Field Approximation
LIS	Local Interstellar Spectrum
NM	Neutron Monitor
NMDB	Neutron Monitor Database
SSN	Sunspot Number
HCS	Heliospheric Current Sheet
HMF / IMF	Heliospheric / Interplanetary Magnetic Field
A	Heliospheric magnetic polarity state (+1 / −1), 22-year Hale cycle
GLE	Ground-Level Enhancement (rare, large solar energetic particle event)
ICME	Interplanetary Coronal Mass Ejection
RMSE / MAE	Root-Mean-Square Error / Mean Absolute Error
POD / FAR	Probability of Detection / False Alarm Ratio
MLP	Multilayer Perceptron
RF / GBR	Random Forest / Gradient Boosting Regressor

Appendix A. Experimental Pipeline Summary (Pseudocode)

The following pseudocode summarizes the two simulation-and-training pipelines used to produce the results in Section 3, to support independent reproduction of the methodology on real NMDB/OMNI data.

```
LONG-TERM MODULATION SIMULATION (Section 2.3, 3.1-3.2):
SIMULATE monthly SSN over N solar cycles using skewed-Gaussian cycle profiles.
DERIVE heliospheric B and HCS tilt angle as SSN-correlated proxies with noise.
ASSIGN polarity state A(t) = +-1, flipping at alternating solar maxima.
COMPUTE modulation potential phi(t) via a polarity-dependent lag/amplitude
  relation to B and tilt angle (drift asymmetry, Section 2.1).
CONVERT phi(t) to GCR intensity (%) via the force-field approximation
  relative to a solar-minimum reference epoch.
ANALYZE SSN-phi and SSN-GCR correlation; plot hysteresis loop by polarity.
SHORT-TERM FORBUSH-DECREASE FORECASTING (Section 2.4-2.6, 3.3-3.8):
SIMULATE hourly solar wind drivers V, N, B, Bz (as in the companion Dst study).
COMPUTE turbulence proxy and INTEGRATE diffusive-barrier ODE for GCR intensity.
BUILD lagged/rolling features; SPLIT chronologically (70/15/15).
FOR EACH horizon h in {6, 24, 72} hours:
    FOR EACH model in {Persistence, Linear, RandomForest, GradientBoosting, MLP}:
      TRAIN, PREDICT, COMPUTE RMSE/MAE/R2/correlation/POD/FAR.
```

EXTRACT Random Forest feature importances (24h horizon model).

ABLATE feature groups; REPEAT MLP training across 3 seeds for robustness.

Both pipelines are directly transferable to real observational data: the synthetic generation steps (1–4 and 7–8) would be replaced by ingestion of real SSN/OMNI/NMDB records, with the remaining feature-engineering, splitting, training, and evaluation steps applied unchanged.

Appendix B. Comparison of Injection/Decay Timescales Across the Two Studies

Table 12 compares the injection/decay structure of the ring-current model used in the companion Dst study with the diffusive-barrier model used here, highlighting the shared mathematical form (a driver-dependent injection term balanced against a linear relaxation term) and the physically distinct timescales and driving variables.

Table 12. Structural comparison of the ring-current (Dst) and diffusive-barrier (GCR/FD) injection/decay models

Physical process	Ring current (Dst)	Diffusive barrier (GCR/FD)
Governing equation form	$dX/dt = Q(t) - X/\tau$	$dX/dt = Q(t) - X/\tau$
Primary driver of Q(t)	Southward IMF (V·Bs)	IMF magnitude & speed (B, V)
Decay constant τ	8 hours	75 hours (~3.1 days)
Typical event amplitude	−30 to −590 nT (Dst)	−2% to −40% (GCR intensity)
Dominant feature (ML)	Dst_lag1 (autoregressive)	CR_lag1 (autoregressive)

This structural parallel is the basis for the methodological claim made in Section 3.9 that the benchmarking framework validated across this pair of studies may generalize to other heliospheric relaxation-type space-weather forecasting problems beyond the two considered here.

Acknowledgements

The authors express their sincere gratitude to the scientific community whose pioneering contributions in heliophysics, cosmic ray physics, astrophysics, plasma physics, and space weather research have laid the foundation for the advancements reviewed in this article. The authors acknowledge the invaluable observational data provided by numerous international space missions and ground-based facilities, including Voyager 1 and 2, ACE, Ulysses, SOHO, STEREO, PAMELA, AMS-02, Parker Solar Probe, Solar Orbiter, Aditya-L1, the Oulu Neutron Monitor, and the Neutron Monitor Database (NMDB), which have significantly advanced the understanding of galactic cosmic ray modulation and heliospheric processes. The authors also appreciate the efforts of researchers and organizations that have developed numerical modeling frameworks and openly shared scientific datasets, thereby facilitating progress in heliophysics and space weather research. Finally, the authors thank their respective institutions for providing the academic environment and research support necessary for the preparation of this review article.

References

[1] L.J. Gleeson, W.I. Axford, Astrophys. J., 1968, 154, 1011–1026.

[2] M.S. Potgieter, Living Rev. Sol. Phys., 2013, 10, 3.

[3] I.G. Usoskin, K. Alanko-Huotari, G.A. Kovaltsov, K. Mursula, J. Geophys. Res., 2005, 110, A12108.

[4] I.G. Usoskin, G.A. Bazilevskaya, G.A. Kovaltsov, J. Geophys. Res., 2011, 116, A02104.

[5] S.E. Forbush, J. Geophys. Res., 1954, 59, 525–542.

[6] J.A. Alhassan, O. Okike, A.E. Chukwude, arXiv:2203.08883, 2022.

[7] A. Mishev, I. Usoskin, J. Space Weather Space Clim., 2018, 8, A19.

[8] Author(s), Cosmic-Ray-Constrained LSTM Model for Geomagnetic Storm Prediction, arXiv:2512.22003, 2025.

[9] E.N. Parker, Phys. Rev., 1958, 110, 1445–1449.

[10] R.A. Caballero-Lopez, H. Moraal, J. Geophys. Res., 2004, 109, A01101.

[11] J. Väisänen, K. Mursula, I.G. Usoskin, G.A. Kovaltsov, J. Geophys. Res. Space Physics, 2023, 128, e2022JA031177.

[12] K. Alanko-Huotari, K. Mursula, I.G. Usoskin, G.A. Kovaltsov, J. Geophys. Res., 2007, 112, A08101.

[13] I.G. Usoskin, G.A. Bazilevskaya, G.A. Kovaltsov, J. Geophys. Res., 2017, 122, 3875–3887.

[14] C. Corti, V. Bindi, C. Consolandi, K. Whitman, Astrophys. J., 2016, 829, 8.

[15] J. Gieseler, B. Heber, P. Herbst, J. Geophys. Res. Space Physics, 2017, 122, 10,964–10,976.

[16] X. Song et al., Mon. Not. R. Astron. Soc., 2026, 547, stag320.

[17] J. Engelbrecht, F. Di Felice, Adv. Space Res., 2020, 65, 2338–2353.

[18] M. Laurenza, E.W. Cliver, M. Storini, Astrophys. J., 2019, 873, 106.

[19] Du, Y., Song, X., & Luo, X. (2024). Deep Learning the Forecast of Galactic Cosmic-Ray Spectra. The Astrophysical Journal Letters, 978.

[20] O. Okike, Forbush decreases: Algorithm generated dataset, Data in Brief, Volume 33, 2020, 106463, ISSN 2352-3409, https://doi.org/10.1016/j.dib.2020.106463..

[21] David Pelosi, Fernando BarÃ£o, Bruna Bertucci, Francesco Faldi, Emanuele Fiandrini, Alejandro Reina Conde, Miguel Orcinha, Nicola Tomassetti, A forecasting framework for galactic cosmic ray flux in space weather applications, Advances in Space Research, Volume 76, Issue 9, 2025, Pages 5700-5713, ISSN 0273-1177, https://doi.org/10.1016/j.asr.2025.08.022..

[22] Agarwal, R., Mishra, Rajesh Kumar and Mishra, Divyansh and (2026). Recent Solar Events and High-Energy Cosmic Ray Particles, International Journal of Science,Engineering and Technology, Volume 14, Issue 2, 2026, ISSN (Online): 2348-4098 ISSN (Print): 2395-4752, 1-15.

[23] R.K. Mishra, D. Mishra, R. Agarwal, Real-time space weather prediction and geomagnetic storm forecasting, International Journal of Science, 14, 3, 1-11, 2026.