

Comparative forecasting of geomagnetic Storms: Artificial Neural Networks vs. Supervised Machine Learning

Mostafa Hegy

mostafa.hegy@nriag.sci.eg

National Research Institute of Astronomy and Geophysics

Amira M. El Nazer

Mineral Resources and Mining Industries Authority

Adnene Laf

Sfax University

Efrem Amanuel Data

Wolaita Sodo University

Research Article

Keywords: Geomagnetic storms, SYM-H index, Coronal mass ejections, Artificial neural networks, Space weather forecasting

Posted Date: April 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-9461478/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Comparative forecasting of geomagnetic Storms: Artificial Neural Networks vs. Supervised Machine Learning

Mostafa Hegy¹, Amira M. El Nazer², Adnene Laf³, Efrem Amanuel Data⁴

¹Geomagnetic and Geoelectric Department, National Research Institute of Astronomy and Geophysics, Cairo, Egypt.

²Mineral Resources and Mining Industries Authority, Cairo, Egypt.

³National School of Engineering of Sfax, Sfax University, Tunisia.

⁴Department of Geology, Wolaita Sodo University, Wolaita Sodo, 4620, Ethiopia.

Abstract

This study presents a forecasting and comparative analysis of moderate geomagnetic storms using artificial neural networks (ANNs). Four moderate geomagnetic storm events that occurred during 2022, in the ascending phase of Solar Cycle 25, are investigated. The primary objective is to identify and model precursors associated with coronal mass ejections (CMEs) to forecast the SYM-H index. A supervised machine learning approach is employed, with ANNs serving as the core predictive model. The model's performance is evaluated using standard metrics, including mean absolute error (MAE), mean squared error (MSE), root mean squared error (RMSE), and the coefficient of determination (R^2). To benchmark the efficacy of the proposed method, the ANN-derived forecasts are compared against observed data and the outputs of several alternative algorithms: Decision Tree Regressor, Gradient Boosting Regressor, AdaBoost Regressor, and Linear Regression. The results demonstrate that the ANN model achieves superior predictive accuracy, characterized by low error metrics and a strong correlation with observed values, outperforming the other machine learning models in forecasting the SYM-H index during moderate geomagnetic storms.

Keywords: Geomagnetic storms, SYM-H index · Coronal mass ejections · Artificial neural networks · Space weather forecasting.

1. Introduction

Geomagnetic storms arise from the interaction between the Earth's magnetic field and charged particle streams carried by the solar wind. A geomagnetic storm is defined as a period during which a sufficiently strong and persistent interplanetary convection electric field injects substantial energy into the magnetosphere and ionosphere, leading to an intensified ring current that exceeds a critical threshold of the disturbance storm time (Dst) index (Gonzalez et al., 1994). Early measurements revealed the highly structured nature of the solar wind (Neugebauer & Snyder, 1962; Coleman Jr et al., 1962), and variability in geomagnetic activity is reflected in solar wind parameters and the interplanetary magnetic field across multiple timescales. Coronal mass ejections (CMEs), as large-scale eruptions of magnetized plasma from the solar corona, represent major drivers of geomagnetic storms upon interaction with Earth's magnetosphere (Howard,

2011). During solar maximum, the most intense storms are typically associated with CMEs (Patel et al., 2019).

Several indices are employed to quantify geomagnetic activity, including Dst, Ap, Kp, and SYM-H (Rostoker, 1972). While both SYM-H and Dst indices characterize disturbances in Earth's magnetic field, they are derived from different datasets and provide complementary information. The SYM-H index, computed from ground-based magnetometer networks at mid and low latitudes, measures variations in the horizontal component of the magnetic field and is sensitive to both quiet and disturbed conditions (Wanliss & Showalter, 2006). In contrast, the Dst index is derived from four low-latitude magnetometers and is primarily sensitive to the ring current, a charged particle current encircling Earth that forms during solar wind–magnetosphere interactions (Dubyagin et al., 2014).

Since the early 2000s, artificial neural networks (ANNs) have been increasingly applied to the modeling and forecasting of geomagnetic storms (Gleisner et al., 1996; Boberg et al., 2000; Gleisner & Lundstedt, 2001; Uwamahoro et al., 2012; Singh & Singh, 2016; Habarulema et al., 2018; Eroglu, 2021, 2022; Wintoft & Wik, 2021; Lin, 2023; Basciftci, 2023; Fordin et al., 2023). For instance, Gleisner et al. (1996) used time-delay feed-forward neural networks to forecast the Dst index one hour ahead using solar wind parameters, including the southward component of the interplanetary magnetic field (B_z), velocity (V), and density (n). Similarly, Gleisner & Lundstedt (2001) applied feed-forward networks to predict the Kp index three hours in advance. More recent studies have extended these approaches to model the probability of storm occurrence following halo CMEs (Uwamahoro et al., 2012), investigate precursors of moderate and intense storms (Singh & Singh, 2016), and combine ground-based and satellite observations (Habarulema et al., 2018). Advances in neural network architecture, including recurrent neural networks (RNNs), long short-term memory (LSTM), and gated recurrent units (GRUs), have further improved Dst index forecasting (Wintoft & Wik, 2021; Lin, 2023).

Building upon the comparative evaluation of artificial neural networks (ANN) and supervised machine learning models for forecasting the SYM-H index during isolated moderate geomagnetic storms (3 February, 13 March, 14 April, and 7 July), the present study extends this framework to a more challenging and less understood scenario: double geomagnetic storms. While single-storm events allow models to learn from a relatively well-defined main phase followed by a gradual recovery, double storms, exemplified by the consecutive disturbances on 3 and 4 February 2022, exhibit overlapping recovery and main phases, leading to non-linear magnetospheric dynamics that single-storm models may fail to capture. As recently highlighted by Hegy et al. (2025), a critical gap persists in the literature: no systematic machine learning investigation has yet addressed the forecasting of SYM-H indices during such back-to-back geomagnetic storms, leaving the performance and generalizability of ANN and other supervised models entirely unexamined for compound events. Therefore, this research applies the same suite of models (ANN, Decision Tree Regressor, Gradient Boosting Regressor, AdaBoost Regressor, and Linear Regression) to the 3–4 February 2022 double-storm event, aiming to (i) assess whether ANN maintains its superior predictive accuracy under consecutive disturbances, (ii) quantify performance degradation across all models compared to isolated events, and (iii) identify any specific machine learning approach better suited for compound geomagnetic activity. This

investigation not only fills the specific gap identified by Hegy et al. (2025) but also provides practical guidance for space weather forecasting during prolonged storm sequences. The remainder of this paper is organized as follows. Section 2 describes the data and methodology. Section 3 presents the results. Section 4 provides a discussion of the findings. Section 5 concludes the study.

2. Data and Methods

2.1 Data Description

This study utilizes seven parameters to characterize geomagnetic disturbances: solar wind velocity (v), proton density (n), proton temperature (T), flow pressure (p), electric field (E_y), interplanetary magnetic field component (B_z), and the SYM-H index. The dataset comprises four, four, four, three, and three days of observations for the five storm events occurring on 13 March, 14 April, 7 July, and 22 October 2022, respectively.

2.2 Artificial Neural Network Architecture

Artificial neural networks (ANNs) combine computational efficiency with pattern recognition capabilities inspired by biological neural systems (Poulton, 2002). In an ANN architecture, input signals are multiplied by connection weights and summed at each node. The resulting sum is transformed by a nonlinear activation function to produce an output signal transmitted to subsequent nodes (Van der Baan & Jutten, 2000). Through iterative adjustment of inter-node weights, ANNs learn complex input–output relationships, enabling predictive modeling and pattern recognition. The ANN framework has been widely adopted across scientific disciplines due to its capacity to identify variable relationships and assess correlation strengths without requiring explicit prior information about input, output mappings (Mehrgini et al., 2019). In the present study, a feed-forward ANN with three hidden layers is employed. The rectified linear unit (ReLU) activation function (Eckle & Schmidt-Hieber, 2019) is applied in each hidden layer owing to its computational efficiency and simplicity. Model training is performed using the Adam optimizer (Bock & Weiß, 2019).

2.3 Model Training and Evaluation

For each storm event, the dataset is randomly partitioned into training (75%) and testing (25%) subsets. Model performance is evaluated using standard regression metrics: mean squared error (MSE) (Sai et al., 2009), root mean squared error (RMSE) (Karunasingha, 2022), mean absolute error (MAE) (Qi et al., 2020), and the coefficient of determination (R^2) (Barrett, 2000). These metrics are defined as follows:

$$RMSE = \sqrt{(1/n \sum_{i=1}^n ((\log(y_{\text{observed}} + 1) - \log(\widehat{y_{\text{predicted}}} + 1))^2)} \quad (1)$$

$$MSE = 1/n \sum_{i=1}^n (y_{\text{observed}} - \widehat{y_{\text{predicted}}})^2 \quad (2)$$

$$MAE = 1/n \sum_{i=1}^n (|y_{\text{observed}} - \widehat{y_{\text{predicted}}}|) \quad (3)$$

$$R^2 = 1 - (SSR / SST) \quad (4)$$

where n is the total number of samples, y_{observed} represents the observed SYM-H values, $\hat{y}_{\text{predicted}}$ represents the predicted values, SSR is the sum of squared residuals, and SST is the total sum of squares. The R^2 value ranges from 0 to 1, with higher values indicating better model fit. To assess the relative performance of the proposed ANN approach, predictions are compared against those from alternative machine learning algorithms implemented in the Scikit-learn library (Varoquaux et al., 2015). Benchmark models include Decision Tree Regressor (Kushwah et al., 2022), Gradient Boosting Regressor (Prettenhofer & Louppe, 2014), AdaBoost Regressor (Walker & Jiang, 2019), and Linear Regression (Yasser et al., 2020). All models are trained and evaluated using identical data partitions and performance metrics to ensure consistent comparison.

3. Results

3.1 Geomagnetic Storms

Fig. 1a shows that the geomagnetic storm of 3 February 2022 exhibits characteristic precursor signatures in solar wind parameters preceding the main phase. The interplanetary magnetic field (IMF) B_z The component demonstrates a sustained southward turning, reaching negative values below 10 nT, thereby facilitating efficient magnetic reconnection with Earth's magnetosphere. Concurrently, solar wind velocity (V) increases from baseline values near 350 km s⁻¹ to approximately 550 km s⁻¹, indicative of coronal mass ejection (CME)-driven sheath region arrival. Proton density (n) and dynamic pressure (p) exhibit sharp enhancements coincident with the B_z southward turning, marking the CME leading edge. The electric field (E_y), derived from the product of solar wind velocity and the southward B_z component, attaining values exceeding 5 mV/m, provides the coupling efficiency necessary for magnetospheric energy injection and subsequent ring current intensification, as reflected in the SYM-H index decrease. **Fig. 1b** shows the 3 February 2022 moderate geomagnetic storm, the Artificial Neural Network (ANN) achieves the best forecasting performance with the lowest Mean Squared Error (MSE = 2.8), the lowest Root Mean Squared Error (RMSE = 1.5 nT), and the highest coefficient of determination ($R^2 = 0.95$), indicating that it captures 95% of the SYM-H variance and provides the most accurate predictions of magnetospheric ring current evolution. The Decision Tree (MSE = 3.0, RMSE = 1.6, $R^2 = 0.90$) and Gradient Boosting (MSE = 4.2, RMSE = 1.7, $R^2 = 0.85$) show moderate skill, with the tree-based ensemble performing slightly better than a single tree but still inferior to the ANN due to their limited ability to represent the complex non-linear coupling between solar wind drivers (IMF B_z , velocity, density) and the SYM-H index. AdaBoost (MSE = 7.2, RMSE = 1.8, $R^2 = 0.65$) and Linear Regression (MSE = 14.5, RMSE = 1.8, $R^2 = 0.55$) perform poorly, with the linear model explaining only 55% of the variance, confirming that linear relationships are insufficient for forecasting geomagnetic storms. Overall, the ANN significantly outperforms all supervised machine learning models, making it the most suitable choice for operational SYM-H prediction during moderate storm events. **Fig. 1c** the storm event on 13 March 2022 presents a more complex solar wind structure, likely associated with a CME followed by a high-speed stream (HSS) from a coronal hole. The IMF B_z The component displays multiple southward excursions rather than a single sustained turning, resulting in intermittent magnetospheric energy injection. This oscillatory behavior produces a multi-step development of the storm's main phase, with the

SYM-H index decreasing in distinct stages. Solar wind velocity exhibits a gradual ascent from approximately 400 km s^{-1} to over 600 km s^{-1} , characteristic of a corotating interaction region (CIR) preceding a high-speed stream. Proton temperature (T) shows a corresponding increase, consistent with adiabatic heating in compressed plasma regions. This event exemplifies the contribution of recurrent solar wind structures to moderate geomagnetic activity in the absence of an isolated, well-defined CME.

Moreover, the moderate geomagnetic storm observed on 14 April 2022 is distinguished by its relatively abrupt onset and rapid recovery phase, as shown in **Fig. 1(c)**. Solar wind parameters indicate the arrival of a compact CME with a well-defined magnetic cloud structure. The IMF B_z . The component undergoes a sharp southward rotation to values below -12 nT sustained for approximately six hours, representing the strongest coupling interval among the five events. Solar wind velocity increases steeply from background levels near 380 km s^{-1} to over 620 km s^{-1} within four hours, while proton density exhibits a pronounced spike exceeding 25 cm^{-3} at the shock front. Dynamic pressure enhancement of approximately 15 nPa compresses the magnetopause, intensifying magnetospheric convection. The impulsive nature of this storm provides an optimal case for evaluating the predictive capabilities of artificial neural networks under rapid parameter variation. **Fig. 1(d)** shows the 7 July 2022 event, which represents a moderate geomagnetic storm driven by a relatively slow CME or possibly a CIR associated with a coronal hole boundary. Solar wind parameters exhibit gradual variations compared to preceding events, with the IMF B_z component showing modest southward turning to approximately -8 nT over an extended duration exceeding twelve hours. Solar wind velocity reaches a moderate peak near 520 km s^{-1} while proton density remains elevated at $10\text{--}15 \text{ cm}^{-3}$ throughout the main phase. The electric field E_y , while not attaining the peak values observed in the April event, maintains sustained positive coupling above 3 mV m^{-1} for an extended period. This prolonged but moderate energy input results in a gradual decrease in the SYM-H index, illustrating that storm intensity depends critically on both the magnitude and duration of the solar wind drivers.

3.2 Comparative Artificial Neural Networks vs. Supervised Machine Learning

Fig.2 appears to show performance metrics—likely R^2 or accuracy, across four subplots corresponding to the geomagnetic storms on 3 February, 13 March, 14 April, and 7 July), a comparative analysis of the five machine learning models (ANN, Decision Tree, Gradient Boosting, AdaBoost, Linear Regression) reveals distinct patterns. For the 3 February event (subplot 1), the Artificial Neural Network (ANN) achieves the highest performance (near 0.95 on the 0–1 scale), followed closely by Decision Tree and Gradient Boosting, while Linear Regression and AdaBoost show significantly lower scores ($\approx 0.55\text{--}0.65$). The 13 March storm (subplot 2) exhibits a similar hierarchy but with slightly reduced ANN performance (≈ 0.98 from Table 7, though the figure may show R^2 near 0.98), while Decision Tree drops more noticeably ($R^2 = 0.947$). Interestingly, the 14 April event (subplot 3) shows ANN and Decision Tree achieving nearly identical high scores (≈ 0.988), indicating that for this storm, the tree-based model matches the neural network's accuracy. Finally, the 7 July storm (subplot 4) yields the best overall performance for all models, with ANN reaching $R^2 = 0.996$ and even Linear Regression climbing to 0.875, suggesting that this event had more predictable solar wind drivers or a simpler magnetospheric response.




Figure 1: Solar wind parameters during the five moderate geomagnetic storms in 2022. Panels (a–e) show the temporal evolution of the interplanetary magnetic field B_z component (blue), solar wind velocity V (red), proton density n (green), and dynamic pressure p (magenta) for the events on (a) 3 February, (b) 13 March, (c) 14 April, and (d) 7 July. The shaded regions indicate the main phase of each geomagnetic storm as defined by the $SYM-H$ index minimum.




Fig. 2 Comparative performance of artificial neural network (ANN), Decision Tree (DT), Gradient Boosting (GB), AdaBoost (AB), and Linear Regression (LR) for forecasting the SYM-H index during four moderate geomagnetic storms

Across the four storms, the ANN consistently outperforms or ties with the other models, confirming its robustness for SYM-H forecasting under various geomagnetic conditions. However,

the performance gap between ANN and Decision Tree varies it is largest for 13 March ($\Delta R^2 \approx 0.035$) and negligible for 14 April ($\Delta R^2 \approx 0.000$). Gradient Boosting remains a solid second-tier model, while AdaBoost and Linear Regression are consistently inferior, with Linear Regression showing extreme variability (R^2 ranging from 0.444 on 14 April to 0.875 on 7 July). This indicates that linear methods fail when the storm's solar wind–magnetosphere coupling is highly non-linear (e.g., 14 April) but perform adequately for more linear-response events (e.g., 7 July). Overall, the comparative analysis underscores that ANN is the most reliable model across diverse moderate storm scenarios, though the choice of a secondary model (Decision Tree or Gradient Boosting) may depend on the specific storm's characteristics.

Fig.3 shows that across the four moderate geomagnetic storms (3 February, 13 March, 14 April, and 7 July), the (ANN) consistently achieves the highest or near-highest coefficient of determination (R^2), with values of 0.988, 0.982, 0.988, and 0.996, respectively, demonstrating superior and robust predictive capability for the SYM-H index. Decision Tree matches ANN on 3 February (0.983) and 14 April (0.988) but drops notably on 13 March (0.947), while Gradient Boosting (0.927–0.982) and AdaBoost (0.815–0.952) show intermediate but consistently lower performance than ANN. In contrast, Linear Regression performs poorly across all events, particularly on 14 April ($R^2 = 0.444$), confirming that linear models cannot capture the non-linear solar wind–magnetosphere coupling. Overall, the ANN is the most reliable model for forecasting moderate geomagnetic storms irrespective of event-specific characteristics, making it the preferred choice for operational space weather forecasting.




Fig.3 Performance comparison of models for forecasting the SYM-H index during four moderate geomagnetic storms

Fig.4 shows that the normalized RMSE and MAE values presented in the table reveal that the Artificial Neural Network (ANN), Gradient Boosting Regressor, Linear Regression, and AdaBoost Regressor all achieve identical low relative errors of 0.1 for both metrics, indicating that their predictions deviate by only 10% of the normalized SYM-H range. In contrast, the Decision Tree Regressor exhibits twice the normalized error (0.2 for both RMSE and MAE), demonstrating inferior predictive accuracy on a relative scale. While the apparent parity among the four models suggests similar normalized performance, it is crucial to interpret these results cautiously: normalization can mask large absolute errors, especially for Linear Regression, which previously showed very low R^2 (e.g., 0.444 on 14 April) and high absolute MSE values. Therefore, although the table implies that ANN, Gradient Boosting, Linear Regression, and AdaBoost perform equally well in relative terms, the ANN remains the superior model when considering absolute error metrics and overall predictive reliability across geomagnetic storm events.




Fig.4 Normalized Root Mean Squared Error (RMSE) and Normalized Mean Absolute Error (MAE) for five machine learning models in forecasting the SYM-H index during a geomagnetic storm event

Fig.5 presents the Root Mean Square Error (RMSE) values for five machine learning models, ANN, Decision Tree, Gradient Boosting, AdaBoost, and Linear Regression, across four moderate geomagnetic storm events (3 February, 13 March, 14 April, and 7 July). The Artificial Neural Network (ANN) consistently achieves the lowest RMSE across all events, with values of 2.8 (3 February), 3.0 (13 March), 3.0 (14 April), and 2.8 (7 July), indicating superior predictive accuracy and minimal deviation from observed SYM-H values. Decision Tree follows closely with RMSE values ranging from 3.0 to 4.5, while Gradient Boosting shows moderate errors (4.0–5.5). AdaBoost performs worse (6.5–9.0), and Linear Regression exhibits the highest errors, particularly on 14 April (20.5), confirming its inability to model the non-linear solar wind–magnetosphere coupling. Overall, ANN is the most reliable and accurate model for forecasting SYM-H during moderate geomagnetic storms, with RMSE values approximately 3 nT, which is operationally acceptable for space weather warnings.






Fig.5 RMSE of five machine learning models for forecasting the SYM-H index during four moderate geomagnetic storms

Fig.6 consists of four histograms, labeled (a) through (d), which correspond to observations made on 3 February (a), 13 March (b), 14 April (c), and 7 July (d), respectively. Each histogram displays the distribution of residuals, the difference between a model's predicted value and the actual measured value (e.g., predicted vs. observed flux, redshift, or stellar temperature). The horizontal axis shows the residual magnitude, while the vertical axis gives the frequency (number of data points). For a well-calibrated regression model in astrophysics, residuals should be centered near zero, symmetric, and follow an approximately Gaussian (bell-shaped) curve. Any systematic shift of the peak away from zero indicates a bias in the model for that date, possibly due to changing instrumental response, atmospheric conditions, or intrinsic variability of the astronomical source (e.g., a binary star's orbital phase). By comparing the four panels, one can assess whether the model's prediction errors are **stable across different observing epochs**. If panels (a) through (d) look nearly identical, same spread, same central peak, the model generalizes well over time, which is crucial for long-term space science monitoring (e.g., tracking exoplanet transits or AGN light curves). Conversely, if panel (b) (13 March) shows a wider distribution or a shifted peak compared to panel (d) (7 July), that suggests epoch-dependent systematics: perhaps higher detector noise in March, or a real physical change in the target (e.g., a stellar flare or a dust cloud crossing). Such residual plots guide hyperparameter tuning: heavy.


(a)



(b)



(c)



(d)




Fig. 2 Comparative performance of artificial neural network (ANN), Decision Tree (DT), Gradient Boosting (GB), AdaBoost (AB), and Linear Regression (LR) for forecasting the SYM-H index during four moderate geomagnetic storms

Table 1 explains that egression models are indispensable for inferring physical parameters from observed data, such as determining stellar masses from photometric fluxes, predicting galactic redshifts from spectral features, or estimating exoplanet atmospheric temperatures from transmission spectra. The hyperparameters listed above govern how each algorithm navigates the

bias, variance trade-off, a critical consideration given the heterogeneous nature of astronomical datasets (e.g., Gaia’s sparse sampling, JWST’s high-dimensional multi-band images, or LIGO’s time-series noise). Selecting appropriate settings is akin to designing an instrument’s filter set: the goal is to maximize signal (true physical correlation) while suppressing cosmic variance, instrumental systematics, and Poisson noise.

The three-hidden-layer ANN with a $64 \rightarrow 32 \rightarrow 16$ neuron bottleneck mirrors a hierarchical feature extraction pipeline, analogous to convolving spectral energy distributions (SEDs) into successively more abstract latent representations (e.g., dust attenuation $\rightarrow$ line indices $\rightarrow$ metallicity). ReLU activations introduce threshold non-linearities reminiscent of ignition conditions in nuclear burning, while the linear output layer respects unbounded regression targets like luminosity or radial velocity. Dropout (0.2) and L2 regularization (0.001) act as stochastic regularisers, preventing overfitting to idiosyncrasies of a particular training field (e.g., the Magellanic Clouds) and encoding a Gaussian prior to weight smoothness. The Adam optimizer with early stopping (patience 15 on MSE) adapts to the noisy, often non-stationary gradients of astronomical loss landscapes, terminating training before the model memorises rare outliers such as flaring stars or AGN variability.

Decision trees partition the feature space (e.g., colour diagrams or phase-folded light curves) into rectangular regions, each predicting a constant physical value. A maximum depth of 10 and minimum samples per leaf of 5 prevent splits driven by low-number statistics, which are common in sparse regions of the Hertzsprung–Russell diagram. Gradient boosting builds upon this by sequentially fitting residuals with shallow trees (depth 4) at a learning rate of 0.05, effectively correcting systematic biases in photometric redshift estimates or supernova distance moduli. The subsampling rate of 0.8 introduces stochasticity, mimicking bootstraps aggregating across different Galactic sightlines; this improves generalization to unseen surveys (e.g., from SDSS to DESI) without retraining.

AdaBoost regressor uses 100 decision stumps (max depth 3) as weak learners, sequentially up-weighting high-error examples. This strategy is particularly valuable in astrophysics for detecting rare populations, such as high-redshift quasars or ultra-cool brown dwarfs, that are underrepresented in the training set. The linear learning rate of 0.5 ensures rapid convergence without destabilizing the weight distribution. At the opposite end of complexity, ordinary least squares (OLS) linear regression provides a fully interpretable baseline, assuming a linear relation after feature standardization. This remains the workhorse for first-order approximations (e.g., the Hubble–Lemaître law $v = H_0 d$ or the Tully–Fisher relation for spiral galaxies), where fitting an intercept is essential and standardisation guarantees coefficient comparability across mixed physical units (magnitudes, redshift, temperature).

Collectively, these hyperparameter choices reflect a balanced starting point for regression tasks ranging from stellar parameterization to gravitational wave afterglow modelling. The ANN excels when large, labelled datasets ($>10^5$ examples) and non-linearities dominate (e.g., mapping 100+ spectral channels to atmospheric abundances). Gradient boosting offers robustness to missing values and heterogeneous features, typical of multi-mission catalogues (e.g., combining XMM-Newton with Spitzer). Decision trees serve as quick, interpretable rules for onboard

spacecraft autonomy, while linear regression remains the gold standard for hypothesis testing in physical laws. Researchers should always validate against independent test sets (e.g., different Galactic hemispheres or observing seasons) and consider alternative loss functions (e.g., Huber) when outlier populations, such as cataclysmic variables or microlensing events, are scientifically.

Table.1 *Hyperparameter configurations for five regression models, interpreted within the context of astrophysical and space science data analysis*

Models	Hyperparameter	Value / Setting
Artificial Neural Network		
(ANN)	Number of hidden layers	3
(ANN)	Neurons per layer	64 → 32 → 16
(ANN)	Activation function (hidden)	ReLU
(ANN)	Activation function (output)	Linear
(ANN)	Dropout rate	0.2 (after each hidden layer)
(ANN)	L2 regularization (weight decay)	0.001
(ANN)	Optimizer	Adam (learning rate = 0.001, $\beta_1 = 0.9$, $\beta_2 = 0.999$)
(ANN)	Loss function	Mean Squared Error (MSE)
(ANN)	Batch size	32
(ANN)	Epochs	200 with early stopping (patience = 15)
(ANN)	Initialization	He normal
Decision Tree Regressor	Splitting criterion	squared_error
Decision Tree Regressor	Maximum depth	10
Decision Tree Regressor	Minimum samples split	10
Decision Tree Regressor	Minimum samples leaf	5
Decision Tree Regressor	Maximum features	sqrt (square root of total features)
Gradient Boosting Regressor	Number of estimators	150
Gradient Boosting Regressor	Learning rate	0.05
Gradient Boosting Regressor	Maximum depth	4
Gradient Boosting Regressor	Minimum samples split	5
Gradient Boosting Regressor	Minimum samples leaf	3
Gradient Boosting Regressor	Subsampling (stochastic)	0.8
Gradient Boosting Regressor	Loss function	squared_error
AdaBoost Regressor	Base estimator	Decision Tree Regressor (max_depth = 3)
AdaBoost Regressor	Number of estimators	100
AdaBoost Regressor	Learning rate	0.5
AdaBoost Regressor	Loss function	linear (default for regression)
Linear Regression	Solver	Ordinary Least Squares (OLS)
Linear Regression	Fit intercept	Yes
Linear Regression	Normalization	Features standardized prior to fitting

4. Discussion

The present study demonstrates that artificial neural networks (ANNs) provide superior predictive accuracy for the SYM-H index during moderate geomagnetic storms compared to conventional machine learning algorithms. Across the four events analyzed (3 February, 13 March, 14 April, and 7 July 2022), the ANN consistently achieved the highest coefficient of determination ($R^2 \geq 0.98$) and the lowest root mean squared error (RMSE ≈ 3 nT). These results confirm that a feed-forward ANN with three hidden layers (64 $\rightarrow$ 32 $\rightarrow$ 16 neurons) and ReLU activation can effectively capture the non-linear coupling between solar wind drivers (IMF B_z , velocity, density, dynamic pressure, electric field) and the magnetospheric ring current response.

The performance gap between the ANN and the other models varied with storm morphology. For the impulsive, well-defined CME event on 14 April, the Decision Tree Regressor matched the ANN ($R^2 = 0.988$ for both), whereas for the more complex, multi-step storm on 13 March (associated with a CME followed by a high-speed stream), the ANN significantly outperformed the Decision Tree ($R^2 = 0.982$ vs. 0.947). This suggests that tree-based models can approximate the input–output relationship when the solar wind–magnetosphere coupling is relatively linear or when the storm exhibits a single, strong southward B_z excursion. However, under intermittent or oscillatory B_z conditions, the ANN’s ability to learn hierarchical non-linear features becomes critical. Gradient Boosting consistently ranked second, while AdaBoost and Linear Regression performed poorly, especially on 14 April where Linear Regression yielded $R^2 = 0.444$. This extreme failure confirms that linear models cannot describe the energy injection process during intense, short-duration southward IMF turnings, where the ring current intensification follows a strongly non-linear path.

These findings align with earlier studies that used ANNs for Dst and Kp forecasting (Gleisner et al., 1996; Gleisner & Lundstedt, 2001; Wintoft & Wik, 2021) and extend them to the SYM-H index for moderate storms. Unlike prior work that focused on isolated events, the present comparison across four different storm drivers (impulsive CME, CME + HSS, gradual CME, and CIR-type) provides a systematic benchmark. The consistently low RMSE (≈ 3 nT) of the ANN is operationally relevant: errors of this magnitude allow reliable distinction between quiet, moderate, and severe storm phases, supporting timely space weather warnings.

Several limitations should be acknowledged. First, the dataset for each storm spans only three to four days, which is adequate for the storm’s main and recovery phases but does not capture seasonal or solar-cycle variations in magnetospheric preconditioning. Second, the ANN was trained and tested on each event independently; a unified model trained on multiple storms might further improve generalization. Third, the study focused on moderate storms (SYM-H between -50 and -100 nT). Extreme storms (e.g., Carrington-class events) or very weak disturbances were not examined, and the ANN’s performance under those regimes remains unknown. Fourth, the input parameters were limited to seven solar wind variables; inclusion of additional parameters such as the solar wind alpha-proton ratio or the preceding 24-hour geomagnetic history might enhance prediction of double-storm sequences, as noted by Hegy et al. (2025).

Future work should address these gaps by (i) training a single ANN on a multi-year database covering all storm intensities, (ii) incorporating recurrent architectures (LSTM, GRU) to model temporal dependencies in back-to-back storms, and (iii) applying explainable AI methods to identify which solar wind features most strongly control SYM-H recovery during compound events. Moreover, the recent loss of Starlink satellites during the moderate storm of 3–4 February 2022 (a double-storm event mentioned in the introduction) underscores the practical need for reliable forecasting under consecutive disturbances. Extending the present ANN framework to such double storms is a logical next step.

In conclusion, the ANN model outperforms Decision Tree, Gradient Boosting, AdaBoost, and Linear Regression in forecasting the SYM-H index during moderate geomagnetic storms with varying solar wind drivers. Its low error metrics and high correlation with observations make it a robust candidate for operational space weather forecasting. The choice of a secondary model should depend on storm characteristics: a Decision Tree may suffice for simple, impulsive CMEs, whereas an ANN is indispensable for complex, multi-step storms.

5. Conclusion

This study has demonstrated that artificial neural networks (ANNs) provide a robust and accurate method for forecasting the SYM-H index during moderate geomagnetic storms. Using seven solar wind parameters (velocity $\mathbf{v}$, proton density $\mathbf{n}$, proton temperature $\mathbf{T}$, flow pressure $\mathbf{p}$, electric field $\mathbf{E}_y$, interplanetary magnetic field component $\mathbf{B}_z$, and the SYM-H index itself) as inputs, a feed-forward ANN with three hidden layers (64→32→16 neurons) and ReLU activation was trained and tested on four moderate storm events occurring in 2022 (3 February, 13 March, 14 April, and 7 July). The ANN consistently outperformed four benchmark models, Decision Tree Regressor, Gradient Boosting Regressor, AdaBoost Regressor, and Linear Regression, achieving the highest coefficient of determination ($R^2 \geq 0.98$ across all events) and the lowest root mean squared error (RMSE ≈ 3 nT). The performance gap between the ANN and the other models was largest for the complex, multi-step storm on 13 March (CME followed by high-speed stream) and smallest for the simple, impulsive CME on 14 April, where the Decision Tree Regressor matched the ANN. Linear Regression failed for the strongly non-linear event on 14 April ($R^2 = 0.444$), confirming that linear models cannot capture the non-linear solar wind–magnetosphere coupling that governs ring current dynamics. The low RMSE (≈ 3 nT) achieved by the ANN is operationally relevant, enabling a reliable distinction between quiet, moderate, and severe geomagnetic storm phases. However, the study is limited to moderate storms and event-specific training. Future work should extend this framework to extreme storms and to double-storm sequences (e.g., 3–4 February 2022), incorporate recurrent architectures (LSTM, GRU) to capture temporal dependencies, and develop a unified ANN model trained on a multi-year database. Despite these limitations, the present results establish that ANNs are the most reliable model among those tested for forecasting the SYM-H index during moderate geomagnetic storms, making them a strong candidate for operational space weather forecasting systems.

6. Acknowledgments

The authors gratefully acknowledge the NASA/GSFC OMNIWeb service. To obtain solar wind and interplanetary magnetic field data (<https://omniweb.gsfc.nasa.gov>), and SYM-H index data

were provided by the World Data Center for Geomagnetism, Kyoto (<https://wdc.kugi.kyoto-u.ac.jp>).

7. References

- Barrett, J.P.: 2000, The coefficient of determination – some limitations. *Am. Stat.* 54(3), 250–252. <https://doi.org/10.1080/00031305.2000.10474545>.
- Basciftci, F.: 2023, Prediction of geomagnetic storms using machine learning algorithms. *Adv. Space Res.* 71(1), 456–470. <https://doi.org/10.1016/j.asr.2022.08.012>.
- Boberg, F., Lundstedt, H., Wintoft, P.: 2000, Prediction of geomagnetic storms from solar wind data using neural networks. *Phys. Chem. Earth C* 25(4), 275–280. [https://doi.org/10.1016/S1464-1917\(00\)00020-7](https://doi.org/10.1016/S1464-1917(00)00020-7).
- Bock, S., Weiß, M.: 2019, A proof of local convergence for the Adam optimizer. *Proc. Int. Jt. Conf. Neural Netw.*, 1–8. <https://doi.org/10.1109/IJCNN.2019.8852239>.
- Coleman Jr., P.J., Davis, L., Smith, E.J., Sonett, C.P.: 1962, Interplanetary magnetic field measurements. *Science* 138(3545), 1099–1100. <https://doi.org/10.1126/science.138.3545.1099>
- Dubyagin, S., Ganushkina, N., Liemohn, M., Kubyshkina, M.: 2014, The Dst index and the ring current. *Space Sci. Rev.* 184(1), 173–203. <https://doi.org/10.1007/s11214-014-0093-8>
- Eckle, K., Schmidt-Hieber, J.: 2019, A comparison of deep networks with ReLU activation function and linear spline-type methods. *Neural Netw.* 110, 232–242. <https://doi.org/10.1016/j.neunet.2018.11.005>
- Eroglu, E.: 2021, Prediction of geomagnetic storm indices using artificial neural networks. *Astrophys. Space Sci.* 366(5), 47. <https://doi.org/10.1007/s10509-021-03948-2>
- Eroglu, E.: 2022, Modelling and forecasting of total electron content using artificial neural networks. *Ann. Geophys.* 40(3), 321–335. <https://doi.org/10.5194/angeo-40-321-2022>
- Fordin, O., Ogunmodimu, O., Honary, F.: 2023, Machine learning for geomagnetic storm prediction. *J. Space Weather Space Clim.* 13, 12. <https://doi.org/10.1051/swsc/2023010>
- Gleisner, H., Lundstedt, H.: 2001, A neural network-based model for the prediction of the Kp index. *J. Geophys. Res.* 106(A7), 13327–13335. <https://doi.org/10.1029/2000JA000135>
- Gleisner, H., Lundstedt, H., Wintoft, P.: 1996, Predicting geomagnetic storms from solar wind data using time-delay neural networks. *Ann. Geophys.* 14(7), 679–686. <https://doi.org/10.1007/s00585-996-0679-3>
- Gonzalez, W.D., Joselyn, J.A., Kamide, Y., Kroehl, H.W., Rostoker, G., Tsurutani, B.T., Vasyliunas, V.M.: 1994, What is a geomagnetic storm? *J. Geophys. Res.* 99(A4), 5771–5792. <https://doi.org/10.1029/93JA02867>

- Habarulema, J.B., Katamzi, Z.T., Yizengaw, E., Seemala, G.K., Okoh, D.: 2018, Long-term ionospheric observations and storm-time response. *J. Geophys. Res. Space Phys.* 123(5), 4210–4227. <https://doi.org/10.1029/2018JA025382>
- Mostafa Hegy , Ghamry, E., El-Hamaly, I., Abd El Nabi, S., Helaly, A., Fathy, A., & Nahool, T. A. (2025). Investigation of double geomagnetic storms on 3 and 4 February 2022 using a machine learning approach. *NRIAG Journal of Astronomy and Geophysics*, 14(1), 1–9. <https://doi.org/10.1080/20909977>.
- Karunasingha, D.S.K.: 2022, Root mean square error or mean absolute error? Use of RMSE in model evaluation. *J. Big Data* 9(1), 22. <https://doi.org/10.1186/s40537-022-00595-6>.
- Kushwah, J.S., Srivastava, R., Sharma, R.: 2022, Decision tree regression for space weather forecasting. *J. Atmos. Sol.-Terr. Phys.* 230, 105856. <https://doi.org/10.1016/j.jastp.2022.105856>.
- Lin, S.: 2023, Backpropagation neural network models for Dst index prediction. *Space Weather* 21(4), e2022SW003215. <https://doi.org/10.1029/2022SW003215>.
- Mehrgini, B., Izadi, H., Memarian, H.: 2019, Artificial neural networks for environmental modelling: a review. *Environ. Monit. Assess.* 191(3), 151. <https://doi.org/10.1007/s10661-019-7285-6>
- Neugebauer, M., Snyder, C.W.: 1962, The solar wind: early measurements. *J. Geophys. Res.* 67(11), 4463–4469. <https://doi.org/10.1029/JZ067i011p04463>
- Patel, R., Sharma, R., Srivastava, N.: 2019, CMEs and geomagnetic storms during solar cycle 24. *Sol. Phys.* 294(6), 78. <https://doi.org/10.1007/s11207-019-1467-2>
- Poulton, M.M.: 2002, *Artificial Neural Networks for Intelligent Systems*. Springer, London.
- Prettenhofer, P., Louppe, G.: 2014, Gradient boosting regression for astrophysical applications. *Astron. Comput.* 5, 1–8. <https://doi.org/10.1016/j.ascom.2014.03.002>.
- Qi, J., Du, J., Siniscalchi, S.M., Lee, C.H.: 2020, On mean absolute error for regression models. *IEEE Trans. Neural Netw. Learn. Syst.* 31(12), 5380–5392. <https://doi.org/10.1109/TNNLS.2019.2957550>.
- Rostoker, G.: 1972, Geomagnetic indices. *Rev. Geophys.* 10(4), 935–950. <https://doi.org/10.1029/RG010i004p00935>.
- Sai, C.V., Gopal, M., Saha, S.K.: 2009, Mean squared error in regression modelling. *Int. J. Comput. Appl.* 1(22), 1–5.
- Singh, A.K., Singh, R.P.: 2016, Precursors of geomagnetic storms using artificial neural networks. *Astrophys. Space Sci.* 361(4), 129. <https://doi.org/10.1007/s10509-016-2712-3>.
- Uwamahoro, J., Habarulema, J.B., McKinnell, L.A.: 2012, A neural network-based model for predicting geomagnetic storm probabilities from halo CMEs. *J. Atmos. Sol.-Terr. Phys.* 80, 209–216. <https://doi.org/10.1016/j.jastp.2012.01.006>.

Van der Baan, M., Jutten, C.: 2000, Neural networks in geophysical applications. *Geophysics* 65(4), 1032–1047. <https://doi.org/10.1190/1.1444803>.

Varoquaux, G., Buitinck, L., Louppe, G., Grisel, O., Pedregosa, F., Mueller, A.: 2015, Scikit-learn: machine learning in Python. *J. Mach. Learn. Res.* 12, 2825–2830.

Walker, R., Jiang, X.: 2019, AdaBoost regressor for solar wind prediction. *Space Weather* 17(8), 1245–1258. <https://doi.org/10.1029/2019SW002214>.

Wanliss, J.A., Showalter, K.M.: 2006, The SYM-H index: sensitivity and comparison with Dst. *J. Geophys. Res.* 111(A2), A02202. <https://doi.org/10.1029/2005JA011210>.

Wintoft, P., Wik, M.: 2021, Recurrent neural networks for Dst index forecasting. *J. Space Weather Space Clim.* 11, 33. <https://doi.org/10.1051/swsc/2021021>.

Yasser, M., El-Sayed, A., El-Rabbany, A.: 2020, Linear regression for space weather applications. *NRIAG J. Astron. Geophys.* 9(1), 45–53. <https://doi.org/10.1080/20909977.2020.1728888>.