Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Combining a Geomagnetic Precursor with Dynamic Time Warping for Solar-Cycle Peak and Trajectory Forecasting

Varun Soman

varunsmn@maharajas.ac.in

Maharaja's College, Ernakulam

Amaranathan Sabarinath

Vikram Sarabhai Space Centre

Research Article

Keywords: Solar cycle forecasting, Geomagnetic aa precursor, Dynamic time warping, Analog forecasting, Sunspot number, Solar Cycle 25

Posted Date: September 23rd, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10909546/v1

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

Additional Declarations: No competing interests reported.

Combining a Geomagnetic Precursor with Dynamic Time Warping for Solar-Cycle Peak and Trajectory Forecasting

Varun Soman[1,2] ·
Amaranathan Sabarinath[3]

© **The author(s)** ●●●●

Abstract

Forecasting the solar cycle is particularly difficult near its onset, when only limited information about the new cycle is available. We combine two complementary sources of information. A geomagnetic *aa*-index precursor provides an early estimate of cycle amplitude, while a historical analog model identifies past solar cycles that most closely resemble the developing cycle and uses them to forecast subsequent sunspot evolution. Cycle similarity is quantified using dynamic time warping (DTW), which allows cycles of different lengths to be compared through flexible alignment along the time axis. The geomagnetic and analog peak estimates are then combined so that the estimate with lower historical uncertainty receives greater influence. The resulting forecast trajectory is adjusted to join the latest observation while keeping sunspot numbers non-negative.

We tested the method by reconstructing forecasts for Solar Cycles 21–24 at six stages of each cycle, using only information available up to each stage. In 20 of the 24 cases, the cycle maximum had not yet been observed. For peak amplitude, the combined estimate has a mean absolute error (MAE) of 16.1 smoothed sunspots, compared with 18.6 for an equal-weight combination and 25.1 for the geomagnetic precursor alone.

Matching the developing cycle against complete historical cycles reduces the mean trajectory error from 26.1 to 20.2 smoothed sunspots, compared with first resampling historical cycles to the length of the observed record. Further tests

V. Soman
varunsmn@maharajas.ac.in

A. Sabarinath
a_sabarinath@vssc.gov.in

[1] Department of Mathematics, Maharaja's College Ernakulam, Kerala, India

[2] Research Scholar, University of Kerala, Thiruvananthapuram, Kerala, India

[3] Vikram Sarabhai Space Centre, Thiruvananthapuram, Kerala, India

show that this improvement is mainly associated with preserving the original lengths of the historical cycles during matching, rather than with the choice between approximate and exact DTW calculations. The trajectory adjustment also prevents artificial flat zero-valued tails during the declining phase.
The resulting empirical uncertainty bands are narrower than those from the earlier DTW method, although their coverage varies considerably between cycles, particularly for Solar Cycle 22. They are therefore interpreted as empirical measures of forecast uncertainty rather than as calibrated prediction intervals. The forecasting procedure has been specified for Solar Cycle 26, allowing the method to be tested on a new solar cycle without changing the rules after the outcome is known.

Keywords: Solar cycle forecasting; Geomagnetic *aa* precursor; Dynamic time warping; Analog forecasting; Sunspot number; Solar Cycle 25

1. Introduction

Forecasting the amplitude and evolution of the solar cycle is an important problem in solar physics. Solar activity modulates the near-Earth radiation environment, thermospheric density, satellite drag, and ionospheric conditions, giving cycle prediction both scientific and space-weather relevance (Hathaway 2015; Petrovay 2020). Proposed approaches include empirical precursors, statistical and analog methods, surface-flux-transport and dynamo models, and data-driven techniques (Pesnell 2008; Petrovay 2020; Nandy 2021). No single source of information is equally useful throughout the cycle.

Precursor methods are most valuable near solar minimum, before the new sunspot cycle has developed substantially. The classical Ohl relationship links low geomagnetic activity near minimum to the subsequent cycle maximum (Ohl 1966; Ohl and Ohl 1979); the long *aa* geomagnetic record makes this relationship particularly useful for retrospective studies (Mayaud 1972). Polar-field precursors have a more direct physical link to the large-scale field that seeds the next cycle (Schatten et al. 1978; Svalgaard, Cliver, and Kamide 2005; Wang and Sheeley 2009; Bhowmik and Nandy 2018; Upton and Hathaway 2018), but precursor methods primarily provide information about cycle amplitude rather than the subsequent sunspot trajectory.

Historical analog methods address a complementary problem. Once a new cycle is sufficiently developed, its observed shape can be matched against previous cycles, whose subsequent evolution then serves as a forecast template. Because solar cycles differ in rise time, peak structure, duration, and decline rate, a similarity measure that allows for differences in timing is useful. Dynamic time warping (DTW) provides this by allowing similar features to occur at different stages of the compared cycles (Sakoe and Chiba 1978; Müller 2007).

Soman and Sabarinath (2026) introduced a DTW-based analog method in which historically similar cycles are selected and combined to forecast the future trajectory and its empirical uncertainty. This method becomes more informative as the target cycle develops, but near cycle onset only limited sunspot information is available. The present work therefore investigates whether a geomagnetic

aa precursor, available from cycle onset, can provide early information about cycle amplitude and complement the analog forecast as the cycle develops.

We also reconsider how historical cycles should be matched to a partially observed target cycle. In Soman and Sabarinath (2026), each complete historical cycle was first resampled to the length of the observed part of the target cycle before DTW matching. Here the historical cycles are instead kept at their original lengths during matching. This direct unequal-length comparison is used as the primary method, while alternative matching choices are examined as robustness checks.

A further practical issue arises when the historical analog trajectory is adjusted to follow the currently observed cycle. A simple vertical shift used to join the forecast to the latest observation can produce negative values during the declining phase. We therefore modify the trajectory so that it remains non-negative.

The method separates estimation of the cycle maximum from forecasting the future trajectory. The *aa* precursor and the analog method each provide an estimate of cycle amplitude. These estimates are combined so that the one with lower historical uncertainty has greater influence. Historical cycles selected using DTW provide the trajectory shape, which is adjusted to join the latest observation while remaining non-negative. We test the method by reconstructing forecasts for Solar Cycles 21–24 at several stages of each cycle, using only earlier cycles and information available up to each forecast stage. Solar Cycle 25 is considered separately as a contemporary example, while Solar Cycle 26 will provide the first test of the complete method on a new solar cycle, with the forecasting rules fixed in advance.

2. Data

2.1. Sunspot Number

We use the monthly 13-month smoothed total sunspot number from Version 2.0 of the International Sunspot Number series, produced by the World Data Center SILSO at the Royal Observatory of Belgium (Clette and Lefèvre 2015, 2016). The 13-month smoothing reduces short-term variability while retaining the main solar-cycle evolution (Hathaway 2015). The series was retrieved in July 2026 and truncated at December 2025, giving 3318 monthly values from July 1749 through December 2025. Recent values may subsequently be revised by the data provider.

Because the 13-month smoother is centered, a smoothed value at month t depends on observations extending approximately six months beyond t. A cutoff defined in terms of the smoothed series therefore does not correspond exactly to the calendar date at which all contributing observations would have been available.

For the primary hindcast analysis, a nominal cutoff T denotes the number of smoothed sunspot values supplied to the forecasting method, so the observed prefix is

$$y(0), \ldots, y(T-1),$$

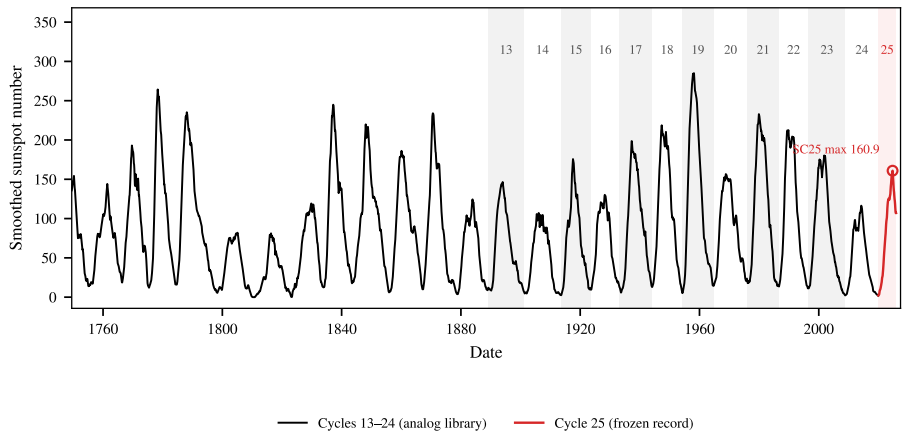


Figure 1. Frozen 13-month smoothed total sunspot number (Clette and Lefèvre 2015) from July 1749 through December 2025, with the adopted solar-cycle segmentation. Cycles 13–24 form the historical analog library, Cycles 21–24 form the principal walk-forward validation set, and Solar Cycle 25 is analyzed separately. The largest observed smoothed value in the frozen Cycle 25 record is 160.9, at elapsed month 57.

with $T = 0$ corresponding to an empty prefix. This gives all competing methods the same observed smoothed prefix and allows direct comparison with Soman and Sabarinath (2026).

We also examine the effect of the approximate six-month smoothing delay. For a forecast issued at elapsed month T, the operational convention is

$$T_{\mathrm{SSN}} = \max(0,\, T - 6), \qquad T_{\mathrm{aa}} = T, \tag{1}$$

where T_{SSN} is the available smoothed-sunspot prefix length and T_{aa} is the geomagnetic cutoff. A more conservative case in which both inputs are lagged by six months is also examined.

Adjacent smoothed values are strongly dependent because consecutive 13-month windows share twelve underlying months. This dependence is taken into account when interpreting forecast errors and uncertainty coverage.

Figure 1 shows the sunspot series and the adopted cycle segmentation.

2.2. Solar-Cycle Segmentation and Month Indexing

We retain the cycle boundaries used by Soman and Sabarinath (2026), listed in Table 1. These boundaries broadly follow minima of the smoothed sunspot series. Keeping the same segmentation allows the present results to be compared directly with the earlier analysis without introducing changes due only to revised cycle boundaries.

Solar Cycles 13–24 form the historical cycle library, and Cycles 21–24 form the principal walk-forward validation set. Solar Cycle 25 is considered separately in Section 5. Within each cycle, elapsed months are indexed from zero, with month 0 denoting the adopted onset. The frozen Cycle 25 record extends from

January 2020 through December 2025 and contains 72 smoothed monthly values; its largest observed value is 160.9 at elapsed month 57.

The adopted onset dates are retrospective cycle boundaries and need not be known exactly at a real forecast origin. We therefore examine the effect of shifting the onset date used for the geomagnetic precursor by ±3, ±6, and ±12 months.

Table 1. Adopted solar-cycle segmentation and realized maxima of the 13-month smoothed sunspot number. Maximum months are elapsed months from the adopted onset (month 0). Cycles 13–24 form the historical cycle library and Cycles 21–24 form the principal validation set. The Cycle 25 entry gives the frozen record through December 2025 and is incomplete; its tabulated maximum is the largest smoothed value observed through that date.

Cycle	Onset	End	Length (months)	Max. SSN	Max. month (elapsed)
13	1889-03	1901-02	144	146.5	58
14	1901-03	1913-07	149	107.1	59
15	1913-08	1923-07	120	175.7	48
16	1923-08	1933-02	115	130.2	56
17	1933-03	1944-01	131	198.6	49
18	1944-02	1954-03	122	218.7	39
19	1954-04	1964-09	126	285.0	47
20	1964-10	1976-02	137	156.6	49
21	1976-03	1986-08	126	232.9	45
22	1986-09	1996-07	119	212.5	38
23	1996-08	2008-12	149	180.3	63
24	2009-01	2019-12	132	116.4	63
25	2020-01	2025-12	72	160.9	57

2.3. Geomagnetic *aa* Index

The geomagnetic precursor is based on the *aa* index, a long-running measure of planetary geomagnetic activity extending back to the nineteenth century (Mayaud 1972). We use the three-hourly series distributed by the International Service of Geomagnetic Indices (ISGI) and calculate monthly means. The resulting monthly series extends from January 1869 through December 2025. Values through December 2024 are definitive, while those for 2025 are provisional.

The analysis uses the classical ISGI *aa* series. Long-term calibration and secular changes in the geomagnetic field can affect its homogeneity (Lockwood et al. 2018); this limitation is considered when interpreting the precursor results.

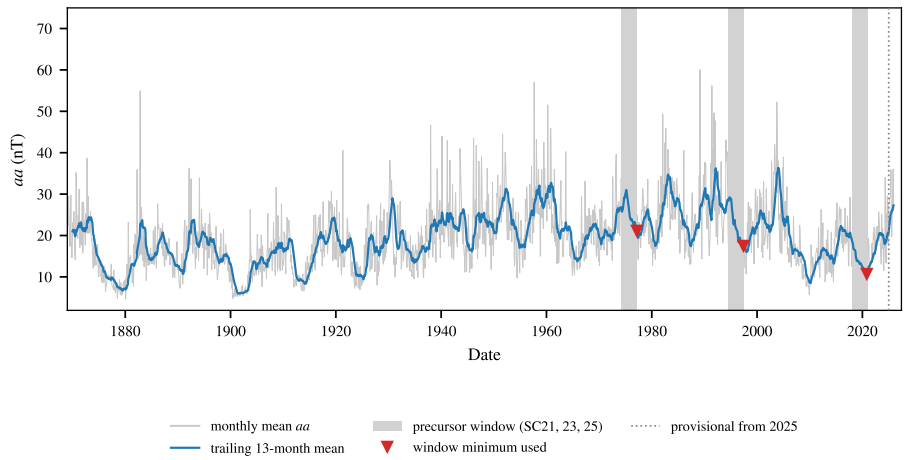


Figure 2. Frozen monthly mean geomagnetic *aa* index (light trace) and trailing 13-month mean from Equation (2) (heavy trace), January 1869 through December 2025. Shaded regions show representative precursor windows for Solar Cycles 21, 23, and 25; downward triangles mark the corresponding minima used by the precursor construction. Values are definitive through December 2024 and provisional during 2025.

To obtain a geomagnetic quantity that is available at the forecast origin, we use a trailing 13-month mean,

$$\bar{a}(t) = \frac{1}{13} \sum_{j=0}^{12} a(t-j), \tag{2}$$

where $a(t)$ is the calendar-month mean *aa* index. Unlike the centered sunspot smoother, $\bar{a}(t)$ depends only on observations at or before month t and is therefore available without a future-data delay.

Figure 2 shows the monthly *aa* series and its trailing 13-month mean.

2.4. Data and Code Availability

All calculations use fixed input data through December 2025. The data, analysis code, executed notebook, output tables, and figures used in this study are available at

https://github.com/varunsmn87/SolarcycleDTWprecursor.

The repository also contains the frozen inputs used to reproduce the Solar Cycle 25 calculations discussed in Section 5.

3. Methods

The forecasting problem is divided into two related tasks: estimating the eventual maximum of the solar cycle and forecasting its future sunspot trajectory. The

geomagnetic *aa* index provides early information about cycle amplitude, while the observed sunspot record is used to identify historically similar cycle shapes.

At a nominal cutoff T, the observed smoothed-sunspot prefix is $y(0), \ldots, y(T-1)$, where $y(t)$ denotes the smoothed sunspot number at elapsed month t. The case $T = 0$ denotes cycle onset, when no target-cycle sunspot trajectory is yet available.

The forecasting procedure consists of four stages:

i) the *aa* precursor provides a predictive distribution for the eventual cycle maximum;
ii) DTW is used to select historical analogs and construct a trajectory shape;
iii) the geomagnetic and analog peak estimates are combined according to their historical uncertainty; and
iv) the analog trajectory is joined to the observed cycle and adjusted toward the combined peak estimate subject to non-negativity.

For every historical target, model fitting and analog selection use only cycles preceding that target. The target cycle contributes only the information admitted at the stated cutoff.

3.1. Geomagnetic Estimate of the Cycle Maximum

The *aa* precursor uses the Ohl relationship between low geomagnetic activity near solar minimum and the amplitude of the following cycle (Ohl 1966). We use the causal trailing 13-month mean $\bar{a}(t)$ defined in Equation (2), so that no future geomagnetic observation is required.

Let c denote the target solar cycle and let $t_{0,c}$ be its adopted onset month. At cutoff T, the geomagnetic predictor $x_c(T)$ is defined as the minimum trailing *aa* value in a window beginning 24 months before the adopted onset:

$$x_c(T) = \min_{t \in [t_{0,c}-24,\, t_{0,c}+\min(T,12)]} \bar{a}(t). \tag{3}$$

At $T = 0$ the window ends at $t_{0,c}$. During the first 12 months of the cycle its right boundary advances as new geomagnetic observations become available; after $T = 12$ the window remains fixed. Thus $x_c(T)$ is a forecast-origin quantity rather than a geomagnetic minimum identified retrospectively from the completed cycle.

For a target cycle c, the precursor relation is fitted using all preceding cycles. Let j index one of these historical cycles, let P_j denote its realized smoothed maximum, and let $x_j(T)$ be its geomagnetic predictor constructed using the same cutoff convention. We fit

$$P_j = \beta_0 + \beta_1 x_j(T) + \varepsilon_j, \qquad \varepsilon_j \sim \mathcal{N}(0, \sigma^2), \tag{4}$$

where β_0 and β_1 are the regression coefficients and σ^2 is the residual variance.

The fitted regression gives a predictive distribution for the unknown maximum A of the target cycle:

$$A \mid x_c(T) \sim t_\nu \left(\mu_{\mathrm{p}}(T), s_{\mathrm{p}}(T)\right), \tag{5}$$

where the predictive mean is

$$\mu_{\mathrm{p}}(T) = \widehat{\beta}_0 + \widehat{\beta}_1 x_c(T),$$

and the predictive scale is

$$s_{\mathrm{p}}^2(T) = \widehat{\sigma}^{\,2}\left[1 + \mathbf{x}_*^{\mathrm{T}}(\mathbf{X}^{\mathrm{T}}\mathbf{X})^{-1}\mathbf{x}_*\right]. \tag{6}$$

Here n is the number of preceding cycles used in the regression, $\nu = n - 2$, $\mathbf{x}_* = (1, x_c(T))^{\mathrm{T}}$, and $\mathbf{X}$ is the $n \times 2$ regression design matrix whose rows are $(1, x_j(T))$. The quantity $\widehat{\sigma}$ is the fitted residual standard deviation. The leading 1 in Equation (6) includes the scatter of a new cycle around the fitted regression, rather than only uncertainty in the fitted mean.

The predictive density is evaluated numerically over $0 \leq A \leq 400$ smoothed sunspots and renormalized over this range. The realized maximum of the target cycle is never used in constructing its own precursor estimate.

3.2. Historical Analog Trajectory

3.2.1. Donor Pool and Shape Normalization

Let $y^{(c)}$ denote the complete smoothed sunspot-number sequence of cycle c. At cutoff T, only its observed prefix $y^{(c)}(0), \ldots, y^{(c)}(T-1)$ is available to the forecasting method.

The trajectory component identifies completed historical cycles whose shapes most closely match this observed prefix. Following Soman and Sabarinath (2026), the primary donor pool contains completed preceding cycles of the same even–odd parity as the target. This restriction is motivated by the approximately 22-year Hale magnetic cycle (Hale et al. 1919; Hathaway 2015). For Solar Cycles 21–24 the resulting candidate pools contain 4, 4, 5, and 5 completed cycles, respectively. Because three donors are normally retained, analog selection operates on a small historical pool. The effect of removing the parity restriction is examined as a robustness check.

DTW is used here to compare cycle shape rather than absolute amplitude. Before matching, each sequence $z(t)$ is therefore scaled to the range $[0, 1]$:

$$\widetilde{z}(t) = \frac{z(t) - \min_s z(s)}{\max_s z(s) - \min_s z(s)}. \tag{7}$$

Here $\widetilde{z}(t)$ denotes the normalized sequence. A numerically constant sequence is represented by zeros. Absolute sunspot amplitude is restored only after the historical analogs have been selected.

3.2.2. Unequal-Length DTW Matching

Dynamic time warping allows similar features in two sequences to occur at different elapsed times. Let $u = (u_1, \ldots, u_N)$ and $v = (v_1, \ldots, v_M)$ be two normalized sequences of lengths N and M, respectively. Their cumulative DTW cost satisfies

$$D(i,j) = |u_i - v_j| + \min\left\{D(i-1,j), D(i,j-1), D(i-1,j-1)\right\}, \tag{8}$$

with $D(0,0) = 0$ and infinite cost along the remaining zero boundaries (Sakoe and Chiba 1978; Müller 2007). The DTW distance used in the primary analysis is the terminal cumulative cost

$$d_{\mathrm{DTW}}(u, v) = D(N, M). \tag{9}$$

No path-length normalization or warping-window restriction is imposed in the primary analysis.

For target cycle c and a completed historical donor j, the normalized observed target prefix is compared directly with the normalized complete donor:

$$d_{c,j}^{\mathrm{exact}}(T) = \mathrm{DTW}\left(\widetilde{y}_{0:T-1}^{(c)}, \widetilde{y}^{(j)}\right), \tag{10}$$

where $d_{c,j}^{\mathrm{exact}}(T)$ is the resulting distance. The two sequences may therefore have different lengths. The historical donor is already fully observed, so no future value of the target cycle enters this calculation. The DTW distance is used only to rank historical donors; the warping path is not interpreted as a physical alignment of solar-cycle phases.

The exact unequal-length rule is adopted as the main matching rule because it directly implements the stated unequal-length cumulative-cost formulation, rather than because it gives the lowest retrospective validation error.

Two alternative matching rules are retained for comparison. The numerical procedure associated with Soman and Sabarinath (2026) first resamples each complete historical donor to the target-prefix length and then applies FastDTW (Salvador and Chan 2007). We also consider exact prefix-to-prefix DTW, in which the observed target prefix is compared only with the first T months of each historical donor. These alternatives test the sensitivity of the forecasts to the way the length difference between a partially observed target and a completed historical cycle is handled.

3.2.3. Constructing the Analog Trajectory

For $T > 0$, the $k = 3$ donors with the smallest DTW distances are retained. For selected donor i, let d_i denote its DTW distance from the target prefix and let L_i denote its original cycle length. The donor receives inverse-distance weight

$$w_i = \frac{(d_i + \epsilon_d)^{-1}}{\sum_{\ell=1}^{k} (d_\ell + \epsilon_d)^{-1}}, \qquad k = 3, \qquad \epsilon_d = 10^{-9}, \tag{11}$$

where ϵ_d prevents division by zero. Other values of k are examined as sensitivity checks.

Because the selected donors can have different durations, the generated trajectory length L is based on their weighted mean duration and is not allowed to end before the current cutoff:

$$L = \max\left\{T, \mathrm{round}\left(\sum_{i=1}^{k} w_i L_i\right)\right\}. \tag{12}$$

After donor selection, the complete normalized donor profiles are resampled to this common length and averaged:

$$g_L(t) = \sum_{i=1}^{k} w_i \mathcal{R}_L \left[\widetilde{y}^{(i)}\right](t), \tag{13}$$

where $\mathcal{R}_L$ denotes linear resampling to length L and $g_L(t)$ is the resulting normalized blended shape. This resampling occurs after donor selection and therefore does not enter the primary unequal-length DTW comparison.

For selected donor i, let $a_i = \max_t y^{(i)}(t) - \min_t y^{(i)}(t)$ denote its amplitude range and $b_i = \min_t y^{(i)}(t)$ its minimum value. The blended profile is returned to the sunspot-number scale according to

$$f(t) = \mathrm{clip}_{[0,1]}\left[g_L(t)\right] \sum_{i=1}^{k} w_i a_i + \sum_{i=1}^{k} w_i b_i, \tag{14}$$

where $\mathrm{clip}_{[0,1]}$ replaces values below zero by zero and values above one by one. The quantity $f(t)$ is the resulting analog trajectory before it is joined to the observed target cycle.

At $T = 0$ there is no observed target shape for DTW matching. Following Soman and Sabarinath (2026), a half-sine cycle-shaped reference is used and two historical donors are retained. Onset forecasts therefore contain no target-specific DTW information and are not used to distinguish between the alternative matching rules.

3.3. Joining the Analog Trajectory to the Observed Cycle

For $T > 0$, the analog trajectory $f(t)$ must join the target cycle at its latest observed value $y(T-1)$. A simple vertical shift gives the additive anchor

$$f_{\mathrm{add}}(t) = f(t) + y(T-1) - f(T-1).$$

This preserves differences within the analog template but can shift part of the declining tail below zero.

A multiplicative alternative is

$$f_{\mathrm{mult}}(t) = f(t)\frac{y(T-1)}{f(T-1)}.$$

For a non-negative template this preserves non-negativity, but it is unstable if $f(T-1)$ is very close to zero.

We therefore use a deterministic rule that retains the additive form whenever it already gives a non-negative future trajectory and otherwise uses the multiplicative form when its denominator is safely positive:

$$f^{\mathrm{anc}}(t) = \begin{cases} f_{\mathrm{add}}(t), & \min\limits_{u \geq T} f_{\mathrm{add}}(u) \geq 0, \\ f_{\mathrm{mult}}(t), & \min\limits_{u \geq T} f_{\mathrm{add}}(u) < 0 \quad \text{and} \quad f(T-1) > \eta, \\ \max\{0, f_{\mathrm{add}}(t)\}, & \text{otherwise}, \end{cases} \tag{15}$$

where $\eta = 10^{-8}$ is a numerical safeguard against division by a value close to zero. The quantity $f^{\mathrm{anc}}(t)$ denotes the anchored analog trajectory.

At $T = 0$ no observed target value is available, so no anchoring operation is applied. For subsequent notation we set $f^{\mathrm{anc}}(t) = f(t)$ at onset.

The analog peak estimate used in the amplitude-combination model is denoted by $\mu_{\mathrm{an}}(T)$. It is defined as the maximum of the constraint-aware trajectory obtained using the earlier numerical donor-selection procedure, before the combined-amplitude rescaling described below. This definition is kept fixed across the different trajectory-matching rules so that changes in trajectory accuracy are not confounded with changes in the analog amplitude supplied to the combination model.

The earlier analog benchmark is retained separately. It uses the same earlier donor-selection procedure with additive anchoring and non-negative truncation. This distinction allows the trajectory-matching rules to be compared while keeping the amplitude input to the combined model unchanged.

3.4. Combining the Two Peak Estimates

The geomagnetic precursor and the analog component both estimate the eventual whole-cycle maximum. The geomagnetic estimate is available from cycle onset, whereas the analog estimate generally becomes more informative as the observed sunspot prefix develops. Rather than specifying a fixed month at which one estimate replaces the other, their relative influence is determined from historical forecast errors.

Let $p_{\mathrm{prec}}(A)$ denote the geomagnetic predictive density for the unknown whole-cycle maximum A from Equation (5). The analog estimate $\mu_{\mathrm{an}}(T)$ is represented as a noisy observation of A:

$$p(\mu_{\mathrm{an}} \mid A) = \mathcal{N}\left(\mu_{\mathrm{an}}; A, \sigma^2_{\mathrm{an}}(T)\right), \tag{16}$$

where $\sigma_{\mathrm{an}}(T)$ describes the historical error of the analog peak estimate at cutoff T.

For target cycle c, let $\mathcal{H}_c$ denote the set of preceding cycles for which the corresponding walk-forward analog hindcast can be formed. If P_h is the realized maximum of historical cycle h and $\mu_{\mathrm{an},h}(T)$ is its analog peak estimate at cutoff T, then

$$\sigma_{\mathrm{an}}(T) = \left[\frac{1}{|\mathcal{H}_c|}\sum_{h\in\mathcal{H}_c}\left(P_h - \mu_{\mathrm{an},h}(T)\right)^2\right]^{1/2}. \tag{17}$$

Here $|\mathcal{H}_c|$ denotes the number of historical analog hindcasts available for target c; for Solar Cycles 21–24 it ranges from six to nine.

The two sources of amplitude information are combined through the working posterior

$$p(A \mid \mu_{\mathrm{an}}) \propto p_{\mathrm{prec}}(A)p(\mu_{\mathrm{an}} \mid A). \tag{18}$$

The combined amplitude estimate is the numerical posterior mean

$$\widehat{A}_{\mathrm{EB}}(T) = \int_0^{400} A\, p(A \mid \mu_{\mathrm{an}})\, \mathrm{d}A. \tag{19}$$

This empirical-Bayes construction is used as a practical way of combining the two sources of amplitude information rather than as a complete probabilistic model of solar-cycle physics. In particular, $\sigma_{\mathrm{an}}(T)$ is estimated from a small historical sample and is then treated as fixed.

A reported estimate of the eventual whole-cycle maximum cannot be smaller than a sunspot value already observed. Let $A_{\mathrm{obs}}(T)$ denote the largest value observed by the cutoff, with $A_{\mathrm{obs}}(0) = 0$ and

$$A_{\mathrm{obs}}(T) = \max_{0 \le t < T} y(t), \qquad T > 0.$$

The reported peak estimate is therefore

$$\widehat{A}_{\mathrm{rep}}(T) = \max\left\{\widehat{A}_{\mathrm{EB}}(T), A_{\mathrm{obs}}(T)\right\}. \tag{20}$$

The same logical lower bound is applied to the other peak estimators when their peak errors are compared.

The distinction between $\widehat{A}_{\mathrm{EB}}(T)$ and $\widehat{A}_{\mathrm{rep}}(T)$ matters after a large cycle maximum has already been observed. The reported whole-cycle maximum must respect that observation, but the remaining declining trajectory should not be forced back toward a peak that has already occurred. The trajectory transformation therefore uses $\widehat{A}_{\mathrm{EB}}(T)$ rather than $\widehat{A}_{\mathrm{rep}}(T)$.

3.5. Adjusting the Trajectory to the Combined Amplitude

The DTW analogs determine the shape and approximate duration of the forecast, whereas $\widehat{A}_{\mathrm{EB}}(T)$ provides its desired amplitude. These two components are combined by rescaling the anchored analog trajectory.

For $T > 0$, let

$$\pi_T = y(T-1)$$

denote the latest observed sunspot value, which is used as the rescaling pivot. At $T = 0$ there is no observed target value, so we instead define

$$\pi_0 = \min_t f(t).$$

The future trajectory is rescaled about this pivot according to

$$\widehat{y}(t) = \pi_T + \lambda\left[f^{\mathrm{anc}}(t) - \pi_T\right], \qquad t \ge T, \tag{21}$$

where $\widehat{y}(t)$ denotes the final deterministic forecast and λ is the applied amplitude-scaling factor.

When the future part of the analog trajectory contains a peak above the pivot, the scaling factor that would attain the desired combined amplitude is

$$\lambda_{\mathrm{des}} = \frac{\widehat{A}_{\mathrm{EB}}(T) - \pi_T}{\max\limits_{t \geq T} f^{\mathrm{anc}}(t) - \pi_T}. \tag{22}$$

Here λ_{des} is the desired scaling factor before imposing the non-negativity requirement.

Let

$$m_f = \min_{t \geq T} f^{\mathrm{anc}}(t)$$

be the minimum of the future anchored trajectory. If λ_{des} is too large, rescaling about the pivot can push the declining tail below zero. The applied factor is therefore

$$\lambda = \begin{cases} \min\left[\lambda_{\mathrm{des}}, \dfrac{\pi_T - \epsilon_f}{\pi_T - m_f}\right], & m_f < \pi_T, \\ \lambda_{\mathrm{des}}, & m_f \geq \pi_T, \end{cases} \qquad \epsilon_f = 10^{-6}. \tag{23}$$

The small constant ϵ_f provides a positive numerical floor. When the feasibility bound is active, the minimum future value equals ϵ_f by construction.

No amplitude rescaling is applied when the relevant template peak has already occurred, when the desired amplitude does not exceed the pivot, or when no future peak remains. If the non-negativity bound is reached before the desired amplitude can be attained, the forecast is labelled *feasibility-limited*. In such cases the reported whole-cycle peak estimate and the maximum actually attained by the remaining forecast trajectory are recorded separately.

3.6. Empirical Uncertainty Envelopes

Trajectory uncertainty is estimated separately from the empirical-Bayes peak calculation using errors from preceding walk-forward forecasts.

For a historical hindcast, let

$$r(t) = y(t) - \widehat{y}(t)$$

denote the forecast residual at month t. For each target, residuals from earlier same-parity cycles are grouped into 12 phase bins according to their normalized position within the solar cycle. Residuals are centered within each bin, and neighbouring bins are pooled when fewer than five residuals are available.

The bootstrap spread is allowed to vary with the amplitude of the forecast. Define

$$\rho(t) = \max\left\{0.35, \sqrt{\frac{\max\{\widehat{y}(t), 0\}}{\max_u \widehat{y}(u)}}\right\}, \tag{24}$$

where $\rho(t)$ is an amplitude-dependent scale factor and u indexes the generated forecast trajectory. Following Soman and Sabarinath (2026), the fixed lower bound of 0.35 prevents the simulated spread from collapsing near solar minimum.

Let $q(t)$ denote the phase bin containing month t. A bootstrap realization b is generated as

$$\widehat{y}^{(b)}(t) = \widehat{y}(t) + s\,\rho(t)\,r^{(b)}_{q(t)}, \tag{25}$$

where $r^{(b)}_{q(t)}$ is sampled with replacement from the centered residuals in phase bin $q(t)$, and s is an overall width multiplier. We use $b = 1, \ldots, 5000$ bootstrap realizations with a fixed random seed.

For each held-out validation cycle, cutoff, trajectory pipeline, and nominal coverage level, the multiplier s is selected using the other three validation cycles. Candidate values are $0.50, 0.55, \ldots, 5.00$, and the smallest value attaining the required mean coverage is retained. If no candidate reaches the nominal coverage level, $s = 5.00$ is used.

The 10th–90th and 1st–99th pointwise percentiles of the bootstrap trajectories define the nominal 80% and 98% envelopes, respectively, with negative lower bounds truncated at zero. Because only three other validation cycles are available for this calibration and coverage is not uniform across cycles, these bands are described as *empirical uncertainty envelopes* rather than calibrated prediction intervals.

3.7. Validation and Model Comparisons

3.7.1. Walk-Forward Hindcasting

The principal validation set consists of Solar Cycles 21–24. Hindcasts are constructed at nominal cutoffs $T = 0, 12, 24, 36, 48$, and 60 months. For every target, precursor fitting, historical analog-error estimation, and donor selection use only cycles preceding that target. The target contributes only the sunspot information admitted under the stated cutoff convention.

Peak accuracy is measured by the absolute difference between the forecast peak and the realized whole-cycle maximum. Because a cutoff occurring after the realized maximum already contains strong information about that maximum, peak errors are also summarized separately for cutoffs at which the realized maximum had not yet entered the observed prefix.

The primary trajectory score is the mean absolute error over future months for which both the observed target and generated forecast are available:

$$\mathrm{MAE}_{\mathrm{future}} = \frac{1}{N_T} \sum_{t=T}^{T+N_T-1} \left| y(t) - \widehat{y}(t) \right|, \tag{26}$$

where L_y is the length of the observed target cycle, $L_{\widehat{y}}$ is the length of the generated trajectory, and

$$N_T = \min(L_y, L_{\widehat{y}}) - T$$

is the number of future months common to the two sequences.

Two supplementary trajectory scores are also reported: one over a common future window of at most 48 months and one extending to the end of the observed target cycle, using terminal-value extension when the generated trajectory ends earlier.

3.7.2. Attribution of Forecast Changes

To determine which parts of the revised framework affect trajectory performance, the forecasting procedure is introduced in stages.

Configuration A reproduces the earlier numerical trajectory of Soman and Sabarinath (2026). Configuration B adds the combined peak estimate while retaining the earlier matching procedure and trajectory transformation. Configuration C additionally introduces the constraint-aware anchoring and rescaling rules. Configuration D, the proposed trajectory method, retains the same combined amplitude and constraints but replaces the earlier matching procedure with exact unequal-length DTW.

Exact prefix-to-prefix DTW is retained as the principal alternative matching rule. An equal-weight average of all eligible historical donors provides a baseline without DTW-based donor ranking.

The earlier numerical matching procedure and the primary unequal-length procedure differ in two respects: the historical donor is either resampled before matching or retained at its original length, and DTW is computed either with FastDTW or with the exact recurrence. A 2×2 comparison crosses these two choices. This allows the effect of donor representation to be distinguished from the effect of the numerical DTW implementation.

3.7.3. Robustness Analyses

The main modelling assumptions are examined through additional sensitivity analyses without redefining the main method. These include prefix-to-prefix matching, path-normalized DTW, an alternative target normalization, unrestricted rather than same-parity donor pools, different values of k, perturbations of the adopted onset date, likelihood tempering, and leave-one-analog-out tests. The information-delay conventions of Equation (1) are also examined to assess the effect of the centered sunspot smoother.

The six cutoffs from a single solar cycle are not independent replicates. Paired trajectory differences are therefore summarized at the solar-cycle level rather than treating all 24 cycle–cutoff cases as independent observations. Cycle-level bootstrap intervals and leave-one-cycle-out matching-rule checks are used as descriptive stability diagnostics only. With four principal validation cycles, they are not interpreted as large-sample statistical inference.

4. Results

Walk-forward hindcasts of Solar Cycles 21–24 are evaluated at nominal cutoffs $T = 0, 12, 24, 36, 48$, and 60 months. Unless stated otherwise, trajectory errors

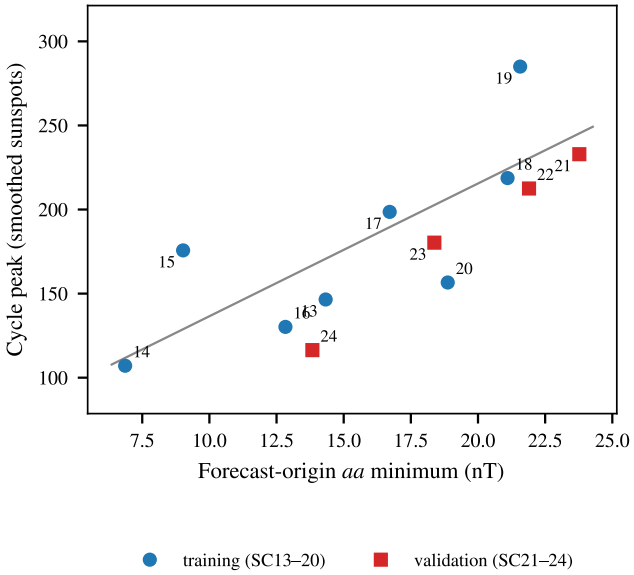


Figure 3. Minimum forecast-origin trailing 13-month mean of the *aa* index available at cycle onset ($T = 0$) against the realized smoothed cycle maximum. Circles denote training Cycles 13–20 and squares denote validation Cycles 21–24. The line is fitted to the training cycles only and is refitted using preceding cycles for each walk-forward target.

are computed over the future period common to the observed cycle and the generated forecast.

4.1. Geomagnetic Information and Peak Forecasts

The forecast-origin geomagnetic predictor is positively associated with the realized smoothed cycle maximum (Figure 3). Using Cycles 13–20 as the initial training sample, the correlation is $r = 0.757$ at cycle onset and $r = 0.782$ when geomagnetic information through the first 12 months is included ($p = 0.030$ and $p = 0.022$, respectively). Given only eight training cycles, these values indicate useful amplitude information rather than a precise estimate of the underlying relation.

Table 2 compares the peak-amplitude forecasts. At cycle onset ($T = 0$), the combined estimate has a mean absolute error (MAE) of 22.07 smoothed sunspots, compared with 23.25 for the geomagnetic precursor and 44.14 for the earlier analog estimate. Because all matching rules use the same onset surrogate at $T = 0$, this result reflects the contribution of the geomagnetic information rather than a difference between DTW matching rules.

The advantage of combining the two sources is most evident early in the cycle. At $T = 12$, the combined MAE is 16.11 smoothed sunspots, compared with 24.36 for the precursor and 39.72 for the earlier analog. At later cutoffs the analog generally becomes more informative, although no single estimator is best at every cutoff.

Table 2. Mean absolute error of the reported whole-cycle peak estimates over Solar Cycles 21–24, in smoothed sunspots. "Earlier analog" denotes the peak from the earlier numerical analog procedure of Soman and Sabarinath (2026); μ_{an} is the analog peak used in the empirical-Bayes combination. The observed-maximum lower bound of Equation (20) is applied to all estimators. "Equal weight" gives equal influence to the precursor and μ_{an}. Bold type denotes the lowest mean error in each row.

T (months)	Precursor	Earlier analog	μ_{an}	Combined	Equal weight
0	23.25	44.14	44.14	**22.07**	26.95
12	24.36	39.72	42.16	**16.11**	18.40
24	24.36	21.08	22.39	14.13	**13.65**
36	24.36	**17.52**	18.92	18.36	20.21
48	24.36	7.25	7.30	**6.79**	10.09
60	24.36	5.32	**4.36**	4.98	11.96
Mean	24.17	22.50	23.21	**13.74**	16.88

Because the observed-maximum lower bound is applied uniformly in Table 2, the earlier-analog values are evaluated under the present reporting convention and are not a verbatim reproduction of the error summary reported in Soman and Sabarinath (2026).

The observed-maximum lower bound affects four of the 24 cycle–cutoff forecasts; the aggregate combined MAE changes from 15.09 before applying the bound to 13.74 after applying it. Restricting the analysis to the 20 cutoffs at which the realized maximum had not yet entered the observed prefix gives a combined MAE of 16.13 smoothed sunspots, compared with 25.11 for the precursor, 26.68 for the earlier analog, 27.53 for μ_{an}, and 18.58 for equal weighting.

The combined estimator is not uniformly better. Using the raw, unprojected amplitudes, it outperforms both component estimates in 11 of the 24 cases. At the earliest cutoffs the precursor tends to overestimate amplitude while the analog tends to underestimate it. Part of the early improvement therefore arises from complementary errors and should not be assumed to persist in future cycles.

As the target cycle develops, the estimated historical error of the analog component generally decreases, so the empirical-Bayes combination gives it greater influence. This is the intended behaviour of the combination rule: the geomagnetic precursor contributes most when little target-specific sunspot information is available, while the analog becomes more influential as the observed trajectory develops.

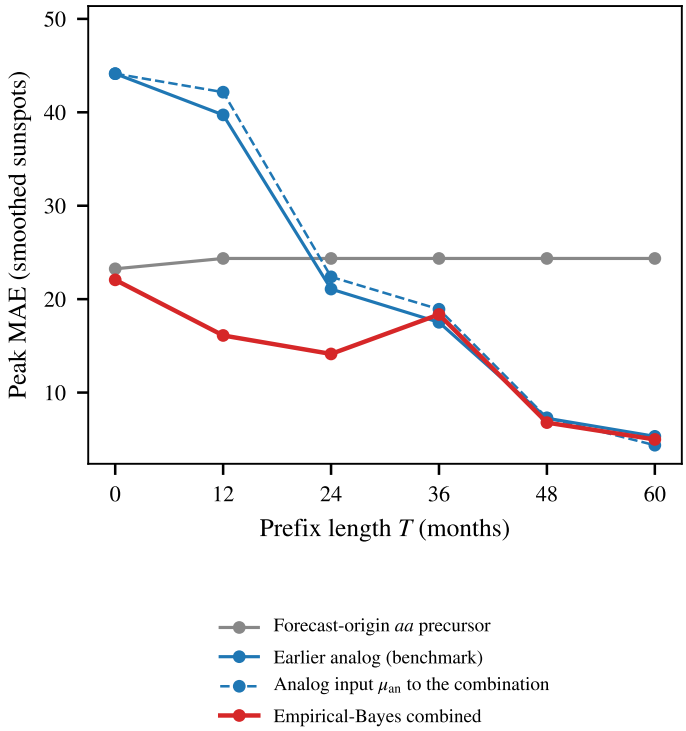


Figure 4. Mean absolute peak error over Solar Cycles 21–24 as a function of nominal cutoff for the geomagnetic precursor, the earlier analog benchmark, the analog estimate μ_{an}, and the empirical-Bayes combined estimate.

4.2. Trajectory Forecasts and DTW Matching

4.2.1. Stepwise Construction of the Forecast

Table 3 shows how the trajectory error changes as the components of the revised method are introduced. The earlier numerical trajectory (Configuration A) has a mean future-trajectory MAE of 29.74 smoothed sunspots. Adding the combined amplitude estimate (Configuration B) gives a modest reduction to 29.07. Introducing the constraint-aware construction (Configuration C) reduces the error further to 26.14. The largest reduction occurs when the matching rule is changed: exact unequal-length DTW (Configuration D) gives a primary mean MAE of 20.20 smoothed sunspots.

The same overall result is seen under the 48-month and full-cycle scoring conventions, so the lower primary error is not explained by differences in generated forecast length. Exact prefix-to-prefix matching also performs well, with a primary MAE of 21.61, but does not improve on the 20.20 obtained with complete unequal-length donors. An equal-weight blend of all eligible donors gives an MAE of 24.82, compared with 20.20 for the exact unequal-length method. This

Table 3. Mean future-trajectory absolute error over Solar Cycles 21–24 under three scoring conventions, in smoothed sunspots. Configurations (A)–(D) introduce the proposed framework stepwise; (E)–(H) examine alternative matching or normalization choices and a baseline without DTW-based donor ranking.

Configuration	Primary	48-month	Full-cycle
(A) Earlier trajectory	29.74	27.58	29.40
(B) Combined peak	29.07	24.86	28.34
(C) Combined peak + constraints	26.14	21.56	25.99
(D) Exact unequal-length	**20.20**	**19.21**	**20.01**
(E) Exact prefix-to-prefix	21.61	21.72	20.97
(F) Path-normalized exact	27.22	21.67	27.07
(G) Alternative normalization	20.38	19.46	20.06
(H) No DTW (equal weights)	24.82	22.25	24.52

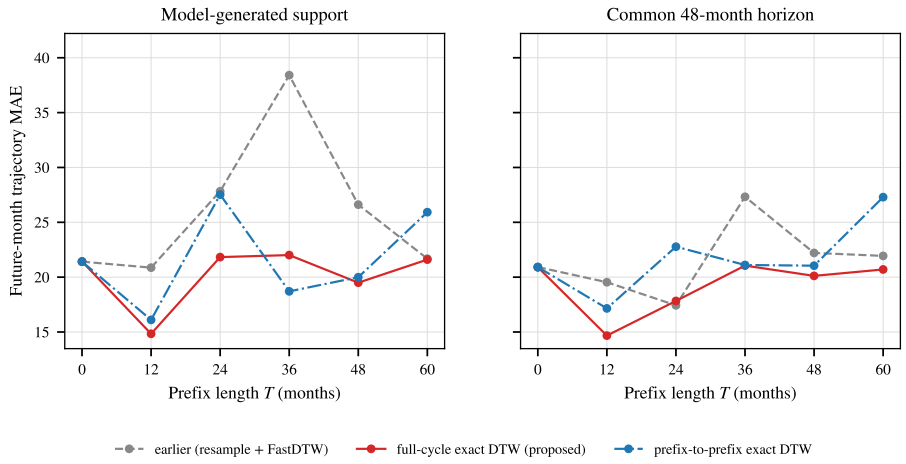


Figure 5. Future-trajectory mean absolute error over Solar Cycles 21–24 for the principal matching rules. The combined amplitude estimate and constraint-aware trajectory transformation are held fixed across matching comparisons. Left: primary model-generated support. Right: common 48-month future window.

suggests that DTW-based donor ranking contributes useful information within the proposed matching formulation, although the small validation set prevents a precise estimate of its advantage.

4.2.2. Why Does the Matching Rule Matter?

The earlier and proposed matching procedures differ in two respects. The earlier numerical procedure resamples each complete historical donor to the target-prefix length before applying FastDTW, whereas the proposed procedure retains the donor at its original length and applies exact DTW.

A 2×2 comparison separates these two changes. With resampled donors, the mean trajectory MAE is 26.136 with FastDTW and 26.195 with exact DTW. With original-length donors, the corresponding values are 20.209 and 20.199. Changing the numerical DTW solver while keeping the donor representation fixed therefore has negligible effect. In contrast, retaining the complete historical donors at their original lengths during matching reduces the mean error by approximately six smoothed sunspots.

The principal difference is therefore associated with donor representation, rather than with the numerical DTW implementation. Exact DTW is retained in the main matching rule because it directly implements the stated cumulative-cost formulation, rather than because greater numerical precision produces the improvement.

The two matching procedures select different donor sets in 12 of the 24 validation cases. When the same donor set is selected, their mean trajectory errors are nearly identical. When the donor set changes, the unequal-length procedure has a lower trajectory MAE by 11.9 smoothed sunspots on average. The improvement is therefore concentrated in cases where the matching convention changes the selected historical analogs, rather than in the subsequent blending of an unchanged donor set.

The improvement is heterogeneous. The mean earlier-minus-unequal-length difference is 5.94 smoothed sunspots, but the median is only 0.28, and unequal-length DTW has the lower error in 14 of the 24 cases. Exact prefix-to-prefix matching is much closer to the primary rule, with a mean paired difference of only 1.41 smoothed sunspots. These results support the unequal-length rule as the main matching rule within the available hindcasts but do not establish it as uniformly superior across solar cycles.

A leave-one-validation-cycle-out check gives a similar qualification. Unequal-length DTW is selected from the remaining three validation cycles in all four folds, but it has the lowest error on only two of the four held-out cycles. The rule is therefore stable as an aggregate choice within the available historical sample, but this does not constitute independent external validation.

4.3. Physical Admissibility of the Forecast Trajectories

The non-negativity rules address a specific artefact of mapping historical analog trajectories onto an observed target. A sufficiently large additive shift can push the declining part of a template below zero, after which simple truncation produces an artificial flat zero tail. The non-negativity adjustment prevents this behaviour and also limits amplitude rescaling when the desired amplification would make the future trajectory negative.

Among the 24 primary validation forecasts, the desired amplitude rescaling is applied in full in 13 cases, is unnecessary in six, and is feasibility-limited in five. In the feasibility-limited cases, the desired combined amplitude cannot be attained without violating the non-negativity constraint. No primary forecast develops an extended exact-zero terminal tail.

The constraint was not introduced to reduce retrospective error. At the 60-month cutoff, in a comparison over Solar Cycles 21–24 using the earlier donor-selection procedure, the constraint-aware construction gives a full-cycle MAE of

Table 4. Peak and trajectory errors under three information-availability conventions. Peak errors are shown at the 12-month issue date and as the mean over all six cutoffs. Trajectory entries are mean future-trajectory MAEs. Absolute trajectory errors across the three information conventions should not be interpreted as identical lead-time experiments because the available smoothed prefix differs. All values are in smoothed sunspots.

	Combined peak MAE		Trajectory MAE		
Information convention	$T = 12$	All cutoffs	Earlier	Exact	Prefix-to-prefix
Nominal prefix	16.11	13.74	26.14	**20.20**	21.61
Sunspot lag only	18.52	15.75	26.91	**20.38**	21.43
Coupled six-month lag	21.73	16.28	27.05	**20.54**	21.44

20.17 smoothed sunspots, compared with 19.40 for simple additive anchoring. Its role is therefore to enforce a physically admissible trajectory, not to optimize historical point accuracy.

4.4. Effect of Information Availability

The primary hindcasts use nominal cutoffs in the centered smoothed sunspot series. Because a centered smoothed value is not fully available until approximately six months after its reference month, we also evaluate the information-delay conventions of Equation (1).

The availability delay has its largest effect on the early peak forecast. At the 12-month issue date, the combined peak MAE increases from 16.11 under the nominal convention to 18.52 with the sunspot lag and to 21.73 with both inputs lagged. The effect is smaller when averaged over all six cutoffs.

The trajectory ranking is unchanged under the two lagged conventions. Exact unequal-length matching has the lowest mean trajectory MAE in all three cases, so the main trajectory result does not depend on treating the centered smoothed sunspot value as available at its nominal reference month.

4.5. Empirical Uncertainty Envelopes

The empirical uncertainty envelopes are narrower for the proposed pipeline than for the earlier DTW pipeline (Table 5). At the nominal 80% level, the mean width decreases from 109.8 to 84.4 smoothed sunspots; at the nominal 98% level, it decreases from 200.5 to 155.9.

The upper calibration limit $s = 5$ was reached in six of the 144 cycle–cutoff–pipeline–level calibrations, all for the earlier pipeline at the nominal 98% level; neither revised trajectory pipeline reached this limit.

The aggregate coverage masks strong cycle-to-cycle differences. For the exact unequal-length pipeline, nominal-80% coverage averaged over cutoffs is 0.958 for Cycle 21, 0.270 for Cycle 22, 0.956 for Cycle 23, and 0.935 for Cycle 24. The corresponding nominal-98% values are 1.000, 0.606, 1.000, and 0.981.

Solar Cycle 22 is a clear failure case. At the 60-month cutoff, only 0.051 of the future months lie inside the nominal-80% envelope and 0.068 inside the

Table 5. Aggregate future-month coverage and mean width of the empirical uncertainty envelopes over Solar Cycles 21–24 and all six cutoffs. Widths are in smoothed sunspots. Nominal levels describe the percentile construction and should not be interpreted as validated coverage probabilities.

Pipeline	Coverage_{80}	Width_{80}	Coverage_{98}	Width_{98}
Earlier DTW	0.767	109.8	0.918	200.5
Exact unequal-length	0.780	84.4	0.897	155.9
Prefix-to-prefix	0.781	81.0	0.902	149.6

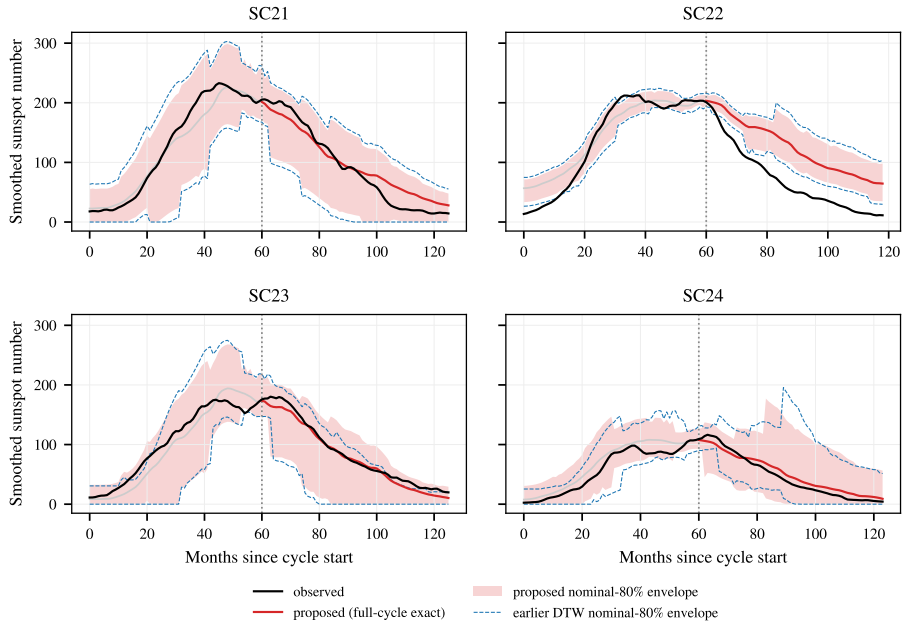


Figure 6. Observed records, proposed point trajectories, and nominal-80% empirical uncertainty envelopes from the exact unequal-length pipeline at a 60-month cutoff (shaded), with the earlier-pipeline nominal-80% bounds shown as dashed lines. The dotted vertical line marks the forecast cutoff.

nominal-98% envelope. The rapidly declining Cycle 22 trajectory is poorly represented by the small historical analog and residual libraries. This interpretation is also supported by the uncertainty calibration: when Cycle 22 is held out, calibration based on the remaining validation cycles selects substantially narrower uncertainty scales than in calibrations that include Cycle 22, indicating that the magnitude of its forecast errors, particularly during decline, is not well represented by the other cycles.

A sensitivity calculation using only residuals from the genuinely future part of each historical hindcast gives essentially the same aggregate result: nominal-80% coverage is 0.777 and mean width is 84.7 smoothed sunspots. The poor

Cycle 22 coverage is therefore not primarily explained by including residuals from already-observed portions of historical cycles.

These results support describing the bands as empirical uncertainty envelopes rather than calibrated prediction intervals.

4.6. Robustness and Stability

The main conclusions are generally stable under the additional checks, although the small historical sample remains an important limitation.

Shifting the adopted onset used to define the geomagnetic precursor window by up to ± 12 months changes the onset forecasts, but combining the precursor with the analog estimate reduces this sensitivity relative to the precursor alone. Tempering the analog likelihood with $\tau = 0.5$ improves the early-cutoff peak errors, particularly at cycle onset, but worsens several later cutoffs and provides essentially no aggregate improvement. There is therefore no clear basis for introducing an additional fitted tempering parameter.

The effect of the same-parity donor restriction depends on the matching rule. For exact unequal-length matching, removing the restriction increases the mean trajectory MAE from 20.20 to 23.49 smoothed sunspots. For prefix-to-prefix matching, it instead decreases slightly from 21.61 to 20.99. The available validation set therefore does not establish a general forecasting advantage for parity restriction.

The ensemble size is also not sharply determined by the available cycles. For the primary score, $k = 3$ gives the lowest error among the tested values, but $k = 2$ and $k = 3$ are nearly indistinguishable under the common 48-month score. Leave-one-donor experiments further show that individual forecasts can change when a single historical analog is removed.

Taken together, the results support three main conclusions. First, combining geomagnetic and analog information is most useful early in the cycle, when the precursor contributes amplitude information not yet available from the sunspot trajectory. Second, the principal reduction in trajectory error is associated with retaining complete historical cycles at their original lengths during DTW matching, which can change the selected historical analogs. Third, the non-negativity adjustment removes this artificial zero-tail behaviour and was not introduced to improve historical forecast accuracy. These conclusions remain subject to the limited number of completed solar cycles available for validation and to the uneven performance of the empirical uncertainty envelopes.

5. Application to Solar Cycle 25 and Forecasting Plan for Solar Cycle 26

Solar Cycle 25 is held outside the principal validation set. The analyses below illustrate how the combined amplitude estimate evolves during the cycle, compare the published trajectory with later smoothed observations, and examine the non-negativity constraint during the declining phase. They are illustrative and do not constitute additional independent validation.

Table 6. Solar Cycle 25 amplitude estimates at successive nominal cutoffs, in smoothed sunspots. "Earlier analog" is the estimate from the earlier numerical donor-selection procedure. "Combined amplitude" is the raw empirical-Bayes estimate before the observed-maximum reporting constraint is applied. The largest value in the frozen Cycle 25 record is 160.9 at elapsed month 57.

T (months)	Earlier analog	Combined amplitude
0	188.7	159.8
12	208.8	172.0
24	184.6	165.6
36	177.5	162.3
48	131.2	133.1
60	173.4	166.8
72	149.5	144.2

5.1. Evolution of the Solar Cycle 25 Amplitude Estimate

The frozen Solar Cycle 25 record contains 72 smoothed monthly values through December 2025. The largest observed value is 160.9 smoothed sunspots at elapsed month 57 (October 2024).

Table 6 shows the analog and combined amplitude estimates at successive nominal cutoffs. Entries through $T = 48$ precede the largest value in the frozen record; those at $T = 60$ and $T = 72$ are post-maximum diagnostics rather than forecasts of an unknown maximum.

At each of the five pre-maximum cutoffs the combined estimate is closer to 160.9 than the earlier analog. Over these five cutoffs, the mean absolute deviation from the frozen maximum is 9.2 smoothed sunspots for the combined estimate and 29.1 for the earlier analog. These figures describe one cycle only and are not included in the validation scores of Section 4.

The sequence also shows the changing balance between the two information sources. The geomagnetic predictor is fixed after $T = 12$, when its prescribed window stops expanding; later changes in the combined estimate reflect only the evolving sunspot prefix and the analog component.

Once the observed prefix contains the realized maximum, the reporting rule of Equation (20) takes effect. For example, the raw combined amplitude at $T = 72$ is 144.2, but the reported whole-cycle peak is 160.9 because that value has already been observed. The raw estimate is retained only as a declining-phase diagnostic.

5.2. The Published Solar Cycle 25 Forecast and Later Observations

The replication calculation reproduces the Solar Cycle 25 forecast of Soman and Sabarinath (2026): the 60-month input prefix selects Solar Cycles 17, 15, and 21 and yields a predicted maximum of 173.4 smoothed sunspots.

We compare this frozen trajectory with the smoothed sunspot values indexed to January–December 2025. The mean absolute error over these twelve months is 18.3 smoothed sunspots. The mean signed error (forecast minus observed) is also +18.3, because the forecast lies above the observed series throughout; both series decline, but the published forecast declines more slowly.

All twelve observed values fall within both the nominal 80% and 98% empirical uncertainty envelopes of the published forecast. This should not be read as twelve independent coverage successes, since adjacent centered 13-month smoothed values share most of their underlying monthly data.

The comparison is not a strict real-time forecasting test. Because the smoothed sunspot number is centered, values near the end of the 60-month input prefix depend partly on later calendar months. It is best understood as an evaluation of the frozen published trajectory against subsequently indexed smoothed values, not as real-time validation of the extensions introduced here.

5.3. Declining-Phase Behaviour and the Non-Negativity Constraint

Solar Cycle 25 illustrates why the non-negativity constraint is needed. During the declining phase, a simple additive shift of the historical analog template can push part of the projected trajectory below zero, and truncating at zero produces an artificial flat tail.

With the earlier donor-selection procedure this problem appears from $T = 36$ onward, producing between 7 and 29 zero-valued future months depending on the cutoff. The same behaviour arises with the primary unequal-length donor selection, for which simple additive anchoring produces between 6 and 28 zero-valued future months. The constraint-aware continuation removes these artificial zero tails in both cases.

The constraint is inactive when the additive continuation is already non-negative and intervenes only when the transformation of the historical template would generate negative values. The Cycle 25 declining phase therefore provides a concrete example of the numerical artefact the rule is designed to prevent.

5.4. Forecasting Plan for Solar Cycle 26

Solar Cycle 26 had not begun at the December 2025 data freeze. The forecast-origin geomagnetic predictor cannot yet be evaluated, and the complete precursor–analog method does not currently provide a Solar Cycle 26 forecast.

For continuity with Soman and Sabarinath (2026), the earlier analog-only baseline can be reproduced. Solar Cycles 18 and 22 give a predicted maximum of 143.3 smoothed sunspots at elapsed month 41. Applying the empirical uncertainty construction of the present study to this baseline gives a nominal 80% range of 110.2–191.0 smoothed sunspots. This differs from the 119.2–173.2 interval reported previously because the two studies use different uncertainty constructions. This calculation is retained only as an analog baseline and is not the Solar Cycle 26 forecast of the integrated method.

The first application of the complete method to Solar Cycle 26 will follow the rules fixed in the present study:

i) Once the onset of Solar Cycle 26 is adopted, the onset date, forecast date, and input-data version will be recorded and retained for that forecast.
ii) The geomagnetic predictor, precursor regression, analog error scale, empirical-Bayes peak combination, unequal-length DTW matching, $k = 3$ inverse-distance weighting, and constraint-aware trajectory construction will be applied as defined in Sections 3.1–3.5, without retrospective retuning.
iii) Completed historical information available at each forecast date will enter the precursor, analog-error, and uncertainty calculations under the same eligibility rules used here. The analog peak μ_{an} in the empirical-Bayes combination will retain the earlier numerical donor-selection definition specified in Section 3.3, so that the peak model used for Solar Cycle 26 is the same one evaluated in the present hindcasts.
iv) Each issued forecast and its corresponding data and code version will be archived before the subsequent outcome is known. Evaluation of these forecasts will use the peak-error, trajectory-error, coverage, and envelope-width measures defined in Section 3.7. Any later methodological change will be reported as a new analysis rather than applied retrospectively to an earlier forecast.

This plan allows the complete precursor–analog method to be tested on Solar Cycle 26 using rules fixed before the cycle outcome is known.

6. Discussion

The method combines two sources of information that become useful at different stages of the solar cycle. Near cycle onset, the geomagnetic *aa* index already provides information about the likely cycle amplitude, whereas the sunspot record is still too short for target-specific analog matching to be strongly informative. As the cycle develops, the growing sunspot trajectory provides progressively more information for historical matching. The method therefore separates amplitude estimation from trajectory construction, allowing each source to contribute where it is most informative.

6.1. Complementary Roles of the Geomagnetic and Analog Information

The walk-forward results support this division of roles. The geomagnetic precursor contributes most near cycle onset, while the analog component generally becomes more informative as the observed sunspot prefix lengthens. The empirical-Bayes combination allows their relative influence to change with the estimated historical uncertainty of the analog component, without requiring a fixed handover month.

The combined estimator is not superior to both components in every individual hindcast. Part of its early advantage arises when the precursor and analog estimates make errors in opposite directions and partly offset one another. The empirical-Bayes step should therefore be understood as a transparent rule for combining two partially complementary sources of amplitude information, rather

than as evidence that they form a complete probabilistic model of solar-cycle amplitude. The available historical record is too short to characterize their joint dependence or the uncertainty in each component's error estimate with high precision.

The choice of the *aa* index is practical rather than exclusive. Polar-field measurements and physics-based dynamo models also provide information about the amplitude of an upcoming cycle (Schatten et al. 1978; Svalgaard, Cliver, and Kamide 2005; Wang and Sheeley 2009; Bhowmik and Nandy 2018; Upton and Hathaway 2018). The long *aa* record is useful here because it allows the precursor and analog components to be examined over a common historical period. The same approach could also incorporate another source of early amplitude information specified in advance.

6.2. Interpretation of the DTW Trajectory Result

The main trajectory finding concerns how the length difference between a partially observed target cycle and completed historical cycles is handled. The earlier numerical implementation first resampled each complete historical donor cycle to the length of the observed target prefix before DTW matching. The primary construction used here instead retains each completed historical cycle at its original length during the similarity comparison.

The validation results indicate that this distinction matters mainly because it can change which historical cycles are selected as analogs. Replacing FastDTW with exact DTW has almost no effect on trajectory error when donor representation is held fixed. The observed reduction in error is instead associated primarily with retaining the completed historical cycles at their original lengths during matching. The result should therefore not be interpreted as evidence that exact DTW is inherently superior to FastDTW for this forecasting problem.

The unequal-length formulation is retained as the main matching rule because it directly implements the cumulative-cost criterion adopted here. Its advantage is not universal: exact prefix-to-prefix matching remains competitive and relative performance varies across validation cycles. An additional caveat is that a raw cumulative DTW cost can reflect the duration of a sequence as well as its shape. The distance metric is therefore used only as an empirical criterion for ranking historical analogs, not as a physical measure of solar-cycle similarity.

A separate issue arises after the analogs have been selected. The resulting historical trajectory must be joined to the observed target and adjusted toward the combined amplitude estimate. A simple additive shift can produce negative sunspot values during cycle decline, while truncation at zero then creates an artificial flat tail. The non-negativity constraint prevents this artefact while leaving trajectories that are already non-negative unchanged. The purpose of this rule is to keep the forecast physically meaningful, rather than to improve historical forecast accuracy.

The present model also does not estimate solar-cycle duration independently. The forecast horizon is inherited from the selected historical analogs. Explicit modelling of cycle duration could be considered separately in future work, but is not required for the forecasting question addressed here.

6.3. Limitations and Uncertainty

The principal limitation is the small number of completed solar cycles available for fitting and validation. This affects the geomagnetic regression, estimation of analog forecast error, donor selection, comparison of matching rules, and construction of uncertainty envelopes. Aggregate performance can therefore conceal substantial dependence on individual historical cycles.

This limitation is especially important for the uncertainty estimates. The revised trajectory methods produce narrower empirical envelopes on average, but the fraction of outcomes falling within the bands varies considerably across cycles. Solar Cycle 22 provides the clearest example: its declining behaviour is poorly represented by the available analog and residual libraries, leading to severe undercoverage. The reported bands should consequently be interpreted as empirical indications of forecast spread rather than as calibrated prediction intervals or joint probability regions for the complete future trajectory.

A further limitation concerns the long-term consistency of the classical *aa* record. Its historical extent is valuable for precursor analysis, but station calibration changes and secular geomagnetic effects can influence its homogeneity (Lockwood et al. 2018). The historical strength of the precursor relationship should therefore not be assumed to persist unchanged in future cycles. Comparison with a homogenized geomagnetic series would be a useful sensitivity study.

The same caution applies to the same-parity restriction on the donor pool. Its motivation comes from the Hale magnetic cycle, but its empirical effect depends on the matching rule. The present validation does not establish same-parity selection as an independent source of forecast skill. It is retained as part of the primary analog construction and its effect is made explicit through the unrestricted-pool sensitivity analysis.

Finally, the historical hindcasts use a common frozen sunspot-data vintage rather than reconstructing the exact SILSO release that would have been available at each historical forecast date. The information cutoffs are enforced within that frozen record, and the lag analysis addresses the delay introduced by the centered smoother, but historical data revisions are not reconstructed. The hindcasts should therefore be regarded as controlled walk-forward experiments under a common data vintage rather than exact recreations of every historical real-time forecast.

6.4. Forecasting Solar Cycle 26

Solar Cycle 25 provides a useful contemporary illustration but is not an independent test of the complete method. Much of the cycle was already known when the present extensions were developed, and the comparison with 2025-indexed values concerns the frozen forecast of Soman and Sabarinath (2026). Its main value here is diagnostic: it shows how the combined amplitude estimate evolves and provides a concrete example of the declining-phase anchoring problem.

The present study tests the method using historical forecasts of completed solar cycles. Solar Cycle 26 will provide the first opportunity to apply the complete method to a new cycle with the forecasting rules fixed beforehand. Once

the cycle begins, the method will use only the information available at each forecast date, without changing the procedure in response to later observations.

The assessment will involve more than whether a single predicted maximum is close to the eventual Cycle 26 peak. It will also examine whether the method continues to provide useful forecasts of cycle amplitude and trajectory as new observations become available, and whether the empirical uncertainty envelopes give a reasonable indication of the forecast errors.

7. Conclusions

This study combines two complementary sources of information for solar-cycle forecasting. The geomagnetic *aa* precursor provides early information about the likely cycle amplitude, while historically similar solar cycles selected using DTW provide information about the subsequent sunspot trajectory. The two peak estimates are combined according to their historical uncertainty, and the forecast trajectory is adjusted so that it remains non-negative.

In historical forecasts of Solar Cycles 21–24, the combined peak estimate has a pre-maximum mean absolute error of 16.1 smoothed sunspots, compared with 25.1 for the geomagnetic precursor alone and 18.6 for an equal-weight combination. For trajectory forecasting, matching the developing cycle directly with complete historical cycles gives a mean future-trajectory error of 20.2 smoothed sunspots, compared with 26.1 when the earlier matching procedure is used with the same amplitude and non-negativity adjustments.

The improvement in trajectory forecasting is associated mainly with keeping the historical cycles at their original lengths during matching, rather than with the numerical choice between FastDTW and exact DTW. The non-negativity rule addresses a separate problem by preventing artificial flat zero-valued tails during the declining phase.

These results are limited by the small number of completed solar cycles available for validation and are not uniform across all cycles. Alternative matching rules remain competitive, and the empirical uncertainty bands show substantial cycle-to-cycle variation. The method should therefore be viewed as a transparent and reproducible forecasting procedure rather than as a universally superior solar-cycle prediction model.

Solar Cycle 26 will provide the first opportunity to test the complete method on a new solar cycle using modelling choices fixed before the outcome is known. This will provide a stronger assessment of its practical forecasting value than further testing on completed historical cycles.

8. Declarations

8.1. Ethics approval

Not applicable. This study uses only publicly available solar and geomagnetic data and does not involve human participants, animals, or personally identifiable information.

8.2. Data and code availability

The sunspot-number data are from SILSO, Royal Observatory of Belgium (https://www.sidc.be/silso) (Clette and Lefèvre 2015, 2016), and the geomagnetic *aa* index is from the International Service of Geomagnetic Indices (https://isgi.unistra.fr) (Mayaud 1972).

The frozen input data, analysis code, executed notebook, figures, and supporting outputs used in this study are publicly available at

https://github.com/varunsmn87/SolarcycleDTWprecursor.

8.3. Author contributions

V.S. conceived the study, developed and implemented the forecasting framework, performed the analyses, prepared the replication materials, and drafted the manuscript. A.S. contributed to the physical interpretation, methodological assessment, and critical revision of the manuscript. Both authors reviewed and approved the final manuscript.

8.4. Funding

The authors received no funding for the submitted work.

8.5. Competing interests

The authors declare no competing interests.

References

Bhowmik, P., Nandy, D.: 2018, Prediction of the strength and timing of sunspot cycle 25 reveal decadal-scale space environmental conditions. *Nature communications* **9**, 5209.

Clette, F., Lefèvre, L.: 2015, SILSO Sunspot Number V2.0, https://doi.org/10.24414/qnza-ac80. Published by WDC SILSO - Royal Observatory of Belgium (ROB). DOI.

Clette, F., Lefèvre, L.: 2016, The new sunspot number: assembling all corrections. *Solar Physics* **291**, 2629.

Hale, G.E., Ellerman, F., Nicholson, S.B., Joy, A.H.: 1919, The magnetic polarity of sun-spots. *The Astrophysical Journal* **49**, 153.

Hathaway, D.H.: 2015, The solar cycle. *Living reviews in solar physics* **12**, 4.

Lockwood, M., Chambodut, A., Barnard, L.A., Owens, M.J., Clarke, E., Mendel, V.: 2018, A homogeneous aa index: 1. Secular variation. *Journal of Space Weather and Space Climate* **8**, A53.

Mayaud, P.-N.: 1972, The aa indices: A 100-year series characterizing the magnetic activity. *Journal of Geophysical Research* **77**, 6870.

Müller, M.: 2007, *Information retrieval for music and motion*, Springer.

Nandy, D.: 2021, Progress in Solar Cycle Predictions: Sunspot Cycles 24–25 in Perspective. *Solar Physics* **296**, 54. DOI.

Ohl, A.I.: 1966, Forecast of sunspot maximum number of cycle 20. *Solnechnye Dannye* **12**, 84.

Ohl, A.I., Ohl, G.I.: 1979, A new method of very long-term prediction of solar activity. In: Donnelly, R.F. (ed.) *NOAA Solar-Terrestrial Predictions Proceedings. Volume 2.* **2**, 258. ADS.

Pesnell, W.D.: 2008, Predictions of solar cycle 24. *Solar Physics* **252**, 209.

Petrovay, K.: 2020, Solar cycle prediction. *Living Reviews in Solar Physics* **17**, 2. DOI.
Sakoe, H., Chiba, S.: 1978, Dynamic programming algorithm optimization for spoken word recognition. *IEEE transactions on acoustics, speech, and signal processing* **26**, 43.
Salvador, S., Chan, P.: 2007, Toward accurate dynamic time warping in linear time and space. *Intelligent data analysis* **11**, 561.
Schatten, K.H., Scherrer, P.H., Svalgaard, L., Wilcox, J.M.: 1978, Using dynamo theory to predict the sunspot number during solar cycle 21. *Geophysical Research Letters* **5**, 411.
Soman, V., Sabarinath, A.: 2026, Analog-Based Forecasting of Solar Cycles Using Dynamic Time Warping and Residual Bootstrapping: Real-Time Application to Cycle 25. *Solar Physics* **301**, 86.
Svalgaard, L., Cliver, E.W., Kamide, Y.: 2005, Sunspot cycle 24: Smallest cycle in 100 years? *Geophysical Research Letters* **32**. DOI. URL.
Upton, L.A., Hathaway, D.H.: 2018, An updated solar cycle 25 prediction with AFT: The modern minimum. *Geophysical Research Letters* **45**, 8091.
Wang, Y.-M., Sheeley, N.R.: 2009, Understanding the Geomagnetic Precursor of the Solar Cycle. *The Astrophysical Journal* **694**, L11. DOI. URL.