I.R. OF IRAN METEOROLOGICAL ORGANIZATION

Journal of the Nivar

Homepage: https://nivar.irimo.ir

Volume 50, Special Issue, Spring and Summer 2026, P. 44-58

(Research Article)

Meteorological-to-surface-soil-moisture propagation timescales across Iran's sub-basins: a national characterization from satellite and reanalysis data (2003–2024)

Hadiani, N[1] | Souri, M[2] | Abolfathi, M[3]

1. Assistant Professor, Department of Geotechnics, Faculty of Civil Engineering, Islamic Azad University, Islamshahr Branch, Islamshahr, Iran.
2. Meteorologist,Department of, Applied Meteorology Development, Tehranm Iran.
3. Head of the Emerging and Knowledge-Based Technologies Department, Tehran.Iran.

Corresponding Author E-mail: n.hadiani@gmail.com

(Received: 20 Jun 2026, Revised: 24 Aug 2026, Accepted: 16 Sep 2026, Published online: 21 Mar 2026)

Abstract

Drought propagates from precipitation deficits into soil-moisture deficits before affecting water supplies, yet whether the timescale of this meteorological-to-surface-soil-moisture transition can be predicted from basin properties has not been established for Iran. Using monthly CHIRPS precipitation and ERA5-Land surface (0–7 cm) soil moisture for 2003–2024 across 66 HydroBASINS level-5 sub-basins, we compute standardized indices (SPI, SSI) and estimate, for each basin, the precipitation accumulation timescale that maximizes the SPI–SSI correlation, with significance corrected for serial autocorrelation. Coupling is strong and ubiquitous: every basin shows a significant relationship ($p < 0.01$), the median correlation is 0.58 (range 0.43–0.74), and 85% of basins exceed $|r| = 0.5$. The surface soil-moisture response is fast, with a median accumulation timescale of two months and 94% of basins responding within two months. The timescale is spatially organized (Moran's I = 0.17, p = 0.002) and is associated in-sample with cropland fraction, vegetation cover, aridity, and snow fraction. Critically, however, these static attributes fail to predict the timescale under spatial cross-validation (gradient-boosted-tree $R^2 = -0.18$, i.e. worse than predicting the mean; 95% CI [−0.19, −0.13]), while neither a location-only model nor a combined model exceeds the no-skill level. A positive-control test confirms the same procedure recovers a genuine attribute signal at this sample size, so the null result reflects an absence of transferable attribute control rather than insufficient data. We conclude that meteorological-to-surface-soil-moisture propagation speed in Iran appears to be set by transient hydroclimatic dynamics and regional spatial structure rather than to be predictable from the fixed basin attributes examined here; the per-basin timescale moreover shows limited temporal reproducibility. These results are a reproducible caution for attribution studies that report in-sample attribute importance without spatial validation. Conclusions are insensitive to the standardization method.

Keywords: drought propagation, Iran, spatial cross-validation, standardized precipitation index, surface soil moisture

1. Introduction

Drought is not a single phenomenon but a cascade. A persistent precipitation deficit (meteorological drought) depletes soil moisture (agricultural drought) and, where deficits persist, propagates into streamflow and storage (hydrological drought) (Van Loon, 2015). The rate at which a precipitation anomaly is transmitted to the soil-moisture store sets how quickly agricultural impacts emerge after the onset of dry conditions, and is therefore central to drought early warning. In arid and semi-arid regions such as Iran, where water resources are already stressed and in-situ monitoring is sparse, an observation-based account of this response timing is valuable but largely absent at the national scale.

The urgency is heightened by a clear drying tendency. Over our 2003–2024 study period, raw annual precipitation declines significantly

Cite this article: Hadiani, N., Souri, M., & Abolfathi, M. (2026). Meteorological-to-surface-soil-moisture propagation timescales across Iran's sub-basins: a national characterization from satellite and reanalysis data (2003–2024). Journal of the Nivar, Special Issue, 44-58. DOI: https://doi.org/10.30467/nivar.2026.587185.1383

E-mail: (2) nedaa.souri@gmail.com (3) abolfathi.irimo@gmail.com

Publisher: Iran Meteorological Organization
DOI: https://doi.org/10.30467/nivar.2026.587185.1383

Print ISSN: 1735-0565
Online ISSN: 2645-3347

(Mann–Kendall, $p < 0.05$) in 18% of basins and surface soil moisture in 24%, with no basin showing a significant increase, and the strongest declines concentrated in northwestern Iran—consistent with widely documented deterioration of the country's water balance. Standardized drought indices, by construction, remove this slow trend and isolate the relative anomaly structure; our propagation analysis therefore complements, rather than replaces, trend assessment, and the two are reported side by side.

Most propagation studies rely on the standardized precipitation index (SPI; McKee et al., 1993) and an analogous standardized soil-moisture index (SSI), identifying the precipitation accumulation period whose SPI most strongly correlates with the SSI; multiscalar standardized indices of this family are now standard for comparing drought phases across regions (Vicente-Serrano et al., 2010). The accumulation period that maximizes this correlation is interpreted as the characteristic response, or propagation, timescale. Global compilations suggest that meteorological-to-soil-moisture propagation typically occurs over one to three months, with substantial regional variation driven by aridity, soil properties, and vegetation cover (Zhu et al., 2021). In arid and semi-arid settings, recent studies in northwest China (Li and Cai, 2024) and semi-arid India (Choudhari et al., 2025; Kumar and Chu, 2024) confirm short propagation lags of one to four months, consistent with the rapid transmission expected where soils are thin and evaporative demand is high. The SSI has been validated for drought monitoring across diverse hydroclimatic settings, from Mediterranean basins (Escorihuela et al., 2025) to humid southeastern U.S. regions (Farahmand et al., 2021), reinforcing its utility as a cross-regional agricultural-drought indicator. A recurring methodological concern is that monthly hydroclimatic series are strongly autocorrelated, so nominal correlation significance is inflated unless the effective sample size is reduced accordingly (Bretherton et al., 1999).

The wider literature treats propagation timing as a spatially variable hydroclimatic response rather than a fixed catchment trait: response times tend to be shorter in arid, water-limited systems and longer where soils, vegetation, or hydrological storage buffer deficits, and they are further modulated by land cover and human regulation (Van Loon, 2015; Huang et al., 2017; Liu et al., 2024). A parallel, methodological literature in spatial ecology and Earth-system modelling cautions that models of spatially structured data can appear skilful while failing to generalize, because spatial autocorrelation inflates apparent predictive performance; spatially blocked cross-validation and explicit spatial null models are recommended for honest assessment (Roberts et al., 2017; Ploton et al., 2020). In hydrology, recent work has begun incorporating spatial autocorrelation explicitly into prediction frameworks (Xu et al., 2024), and machine-learning drought predictions routinely identify SPI as the dominant predictor while revealing spatially heterogeneous secondary controls (Yao, 2024; Jiang et al., 2023). Nested cross-validation frameworks have also been applied to assess model spatial stability in semi-arid rainfall–runoff regionalization (Bouaicha et al., 2024), underscoring the value of spatially explicit validation in water-scarce regions. These strands motivate our design: characterize propagation timing nationally, then test whether it is genuinely predictable from basin attributes once spatial dependence is controlled.

Here we provide the first national, satellite-and-reanalysis characterization of meteorological-to-agricultural drought propagation for Iran. We pursue three aims: (i) to quantify the strength and timescale of SPI–SSI coupling across Iran's sub-basins with autocorrelation-corrected significance; (ii) to map and describe the spatial variation of the propagation timescale; and (iii) to test rigorously whether static basin attributes can predict that timescale out-of-sample, against explicit spatial and positive-control baselines. The third aim is the conceptual core of the paper: we show that apparent attribute-based controls are an artefact of spatial autocorrelation, and that no model—including one using location alone—exceeds the no-skill level out-of-sample. This reframes propagation speed as an emergent, transient property rather than a fixed basin trait, with direct implications for how attribution studies in data-scarce regions should be validated. Our findings complement recent Iran-specific drought projections that document intensifying drying trends and hotspot formation (Khoorani et al., 2024) and align with regional attribution studies identifying climate change as a dominant driver of recent Middle East drought intensification (World Weather Attribution, 2024), by quantifying the agricultural response timing within which early warning must operate.

2 Research Method

Figure 1 summarizes the complete analysis workflow described in this section: the input datasets and study units (top), index construction and propagation-timescale estimation (middle), and the significance, spatial-diagnostic, and predictability checks that lead to the validated results reported in Section 3 (bottom).

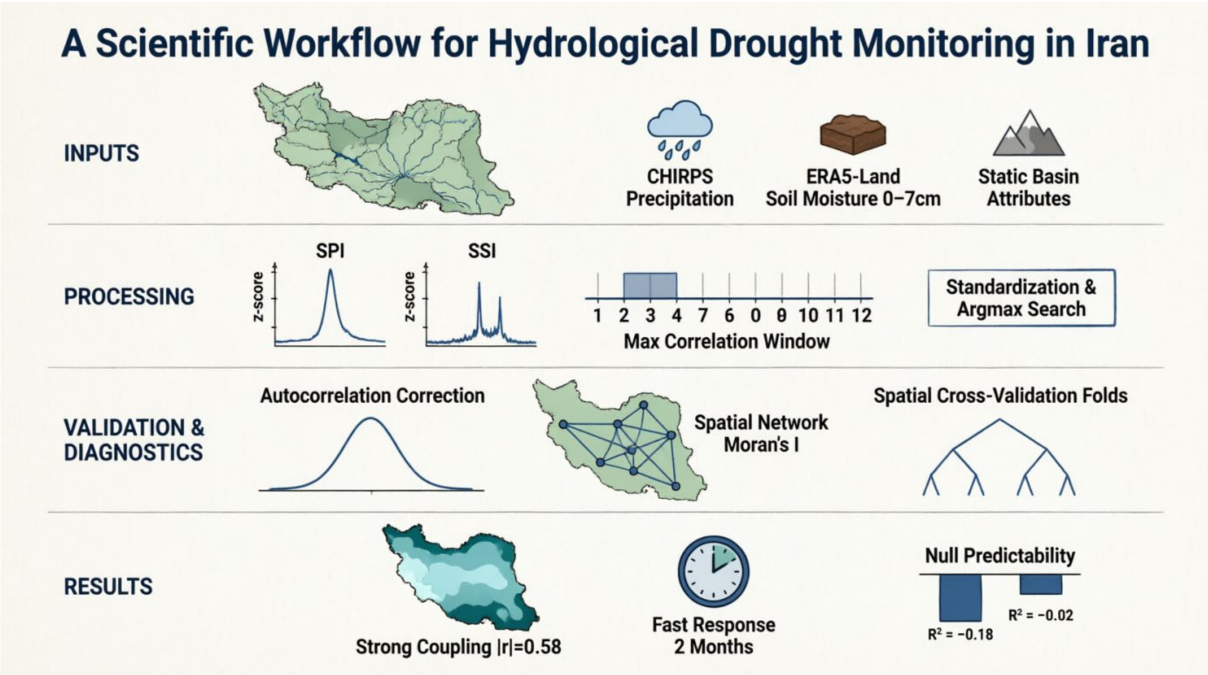

Figure 1. Schematic workflow of the analysis pipeline, from the input datasets through index construction, propagation-timescale estimation, and the significance, spatial-diagnostic, and predictability checks, to the validated results reported in Section 3.

2.1 Study area and analysis units

The study domain is continental Iran, partitioned into 66 HydroBASINS level-5 sub-basins intersecting the national boundary (Lehner and Grill, 2013). Individual sub-basin areas vary considerably, from small headwater catchments to large interior basins, with a mean area of approximately 24,000 km². Basins span a wide hydroclimatic range, from hyper-arid interior plains to snow-influenced mountain headwaters, with a P/PET humidity index (mean annual precipitation divided by potential evapotranspiration) ranging from about 0.01 to 1.10 across the domain (lower values denote more arid conditions). Figure 2 shows the spatial distribution of the 66 sub-basins across the study domain, plotted at basin centroids and coloured by the P/PET aridity gradient.

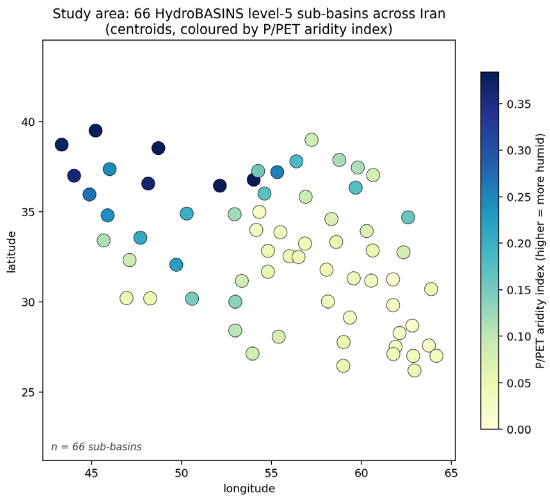

Figure 2. Study area: the P/PET aridity index for each of the 66 HydroBASINS level-5 sub-basins across Iran, plotted at basin centroids.

2.2 Datasets

Monthly precipitation totals were derived from CHIRPS daily precipitation, summed to calendar months (Funk et al., 2015). Soil moisture was taken as the monthly mean of the uppermost ERA5-Land layer (0–7 cm) (Muñoz-Sabater et al., 2021); we therefore characterize the surface soil-moisture response, and interpret the SSI as a surface agricultural-drought proxy rather than a full root-zone measure (see Section 4.2). Static basin attributes—the P/PET humidity index, elevation, slope, long-term mean NDVI, cropland fraction, snow fraction, permanent-water fraction, population density, and an endorheic flag—were computed per basin from SRTM, MODIS, ESA WorldCover, JRC Global Surface Water, and WorldPop. All extraction was performed in Google Earth Engine by reducing each gridded product to basin means for every month from January 2003 to December 2024 (264 months).

We deliberately restrict the analysis to the meteorological-to-agricultural transition. A hydrological stage based on GRACE/GRACE-FO total water storage was evaluated and excluded: the GRACE mascon footprint (on the order of 90,000 km²) greatly exceeds the mean level-5 basin area, so adjacent basins share near-identical storage signals and per-basin estimates are not independent at this resolution.

2.3 Standardized indices

For each basin and variable, a standardized index was computed by removing the seasonal cycle through monthly standardization: values for a given calendar month were converted to z-scores using that month's long-term mean and standard deviation across the record. Formally, for basin i, variable x (precipitation or soil moisture), and calendar month m, the standardized value at time t is given by:

$$z_{i,t} = (x_{i,t} - \mu_i(m)) / \sigma_i(m) \quad (1)$$

where $\mu_i(m)$ and $\sigma_i(m)$ are the calendar-month mean and standard deviation of x_i computed over the full 2003–2024 record. This yields the SPI and SSI on a common, seasonality-free footing.

2.4 Propagation timescale

For each basin we searched accumulation periods of 1–12 months for precipitation and response lags of 0–6 months, and selected the accumulation–lag combination that maximized the absolute Pearson correlation between accumulated SPI and SSI (requiring at least 36 overlapping months). Formally, the selected accumulation a* and lag l* satisfy:

$$(a^*, l^*) = \mathrm{argmax}_{a,l} \, |corr(SPI_{a,t-l}, SSI_t)| \quad (2)$$

over $a \in \{1, \ldots, 12\}$ and $l \in \{0, \ldots, 6\}$, subject to the minimum-overlap constraint above. The accumulation period at this maximum defines the propagation timescale; the lag describes any additional delay.

2.5 Autocorrelation-corrected significance

Because monthly indices are serially correlated, the effective sample size was reduced following the standard lag-1 autocorrelation adjustment (Bretherton et al., 1999), and correlation significance was evaluated against this effective sample size rather than the nominal record length. This correction is applied throughout; nominal (uncorrected) significance is systematically more permissive and is not reported.

2.6 Association and predictability analysis

Associations between the propagation timescale and the static basin attributes were quantified with Spearman rank correlations and are reported as descriptive, in-sample relationships. To test whether these attributes predict the timescale, we trained gradient-boosted regression trees and a linear model and evaluated them under two cross-validation schemes: standard random k-fold and spatial-block k-fold, where folds correspond to geographically contiguous clusters of basins (k-means on basin centroids). The contrast between random and spatial cross-validation isolates the role of spatial autocorrelation in apparent predictability (Roberts et al., 2017). In practice, spatial-block folds were generated by applying k-means clustering to the basin centroid coordinates and using the resulting cluster labels as group identifiers in a grouped k-fold split, which guarantees that all basins in a given spatial cluster are held out together and never appear in both the training and test partitions of the same fold.

2.7 Spatial diagnostics

We quantified spatial autocorrelation in the propagation timescale with Moran's I, using row-standardized six-nearest-neighbour weights based on great-circle distances between basin centroids, and assessed significance with 999 random permutations. The weight matrix followed the standard row-standardized k-nearest-neighbour convention, with each basin connected to its six geographically closest basins and the weights of each basin's connections summing to one, so that

Moran's I reduces to a weighted correlation between each basin's timescale and the average timescale of its nearest neighbours. As an explicit benchmark for the attribute model, we built a purely spatial predictor — referred to hereafter as the location-only, or spatial-neighbour, model — that estimates each held-out basin's timescale as the inverse-distance-weighted mean of the timescales observed at its geographically nearest training-set basins, evaluated under the same spatial-block folds (Ploton et al., 2020). This model uses no physical basin attributes, so it serves as a baseline for testing whether the static descriptors examined here add predictive skill beyond spatial proximity alone. Robustness was tested across four to ten spatial blocks, and sensitivity to the index definition was assessed with a non-parametric empirical-Gaussian standardization in place of the monthly z-score. The identifiability of the timescale was checked by inspecting the median |SPI–SSI| correlation as a function of accumulation window. For context, non-parametric Mann–Kendall trends were computed on raw annual precipitation and soil moisture per basin.

2.8 Power, robustness, and temporal-stability checks

Because predictive conclusions from 66 basins must be interpreted cautiously, we added four checks. First, the spatial-block cross-validated R^2 was recomputed under 50 independent random block partitions, yielding a distribution and 95% interval for each model. Second, as a positive control for statistical power, we generated synthetic targets that are by construction a function of a real attribute (the aridity rank) plus Gaussian noise at weak, moderate, and strong signal-to-noise levels, and verified whether the identical spatial-CV procedure recovers positive skill. Third, we tested a combined model that augments the attribute set with the spatial neighbour estimate, and a dimensionality-reduced model using the first two to four principal components of the standardized attributes. Fourth, we assessed temporal stability by splitting the record into two halves (2003–2013 and 2014–2024), re-estimating the timescale in each, and comparing both the national distribution and the per-basin rank correlation between halves.

3 Measurement, Observation and Calculation

3.1 Ubiquitous and strong meteorological-to-surface-soil-moisture coupling

SPI–SSI coupling is significant in all 66 basins after the autocorrelation correction ($p < 0.01$ in every basin). The median absolute correlation is 0.58 (interquartile range 0.53–0.65; full range 0.43–0.74), and 85% of basins exceed $|r| = 0.5$ (Table 1); all couplings remain significant after Bonferroni correction across the 66 basins (all $p < 0.001$). Coupling strength is highest in the wetter northwest and along the western mountains and lower in parts of the arid interior and southeast, but remains significant everywhere (Figure 3). The transmission of precipitation anomalies into the surface soil-moisture store is therefore a coherent, national-scale feature rather than a property of a few basins.

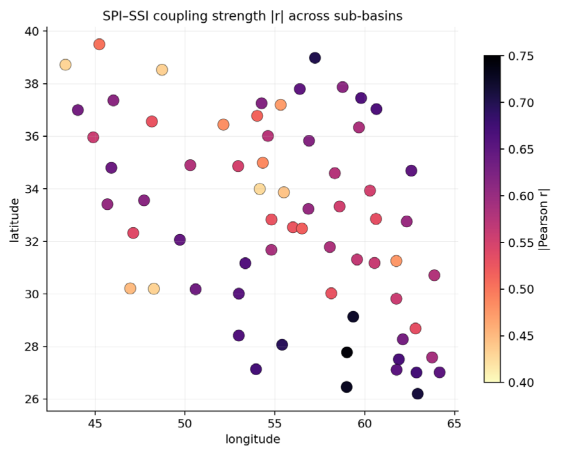

Figure 3. SPI–SSI coupling strength (|Pearson r|) for each sub-basin, plotted at basin centroids. Coupling is significant (p < 0.01) in every basin.

Table 1. Headline characterization of SPI–SSI propagation across the 66 sub-basins.

Metric	Value
Basins (n)	66
Months	264 (Jan 2003–Dec 2024)
Significant coupling ($p < 0.05$ / $p < 0.01$)	100% / 100%
Median |r| (IQR; range)	0.58 (0.53–0.65; 0.43–0.74)
Strong coupling (|r| > 0.5)	85%
Median accumulation timescale (IQR)	2 months (1–2)
Responds within 2 months	94%
Best-fitting response lag	0 months (all basins)

3.2 Overall precipitation and soil-moisture trends over the study period (2003–2024)

Before characterizing relative-anomaly propagation, we first quantify the underlying (raw) precipitation and soil-moisture trends. On raw annual data, precipitation declines significantly in 18% of basins and surface soil moisture in 24%, with no basin exhibiting a significant increase; the median Mann–Kendall τ is negative for both variables (−0.16 and −0.13) (Table 2). The drying is spatially organized, strongest in the northwest and along the western mountains and weakest in the southeast (Figure 4). This unidirectional signal is consistent with widely reported deterioration of Iran's water balance (Khoorani et al., 2024; World Weather Attribution, 2024). Because standardized indices remove this slow component, it does not enter the propagation analysis; we report it as essential context and as motivation for monitoring agricultural-drought response timing under a drying climate.

Table 2. Raw (non-standardized) Mann–Kendall trends over 2003–2024 (context).

Variable	Median τ	Significant drying	Significant wetting
Precipitation	−0.16	18%	0%
Surface soil moisture	−0.13	24%	0%

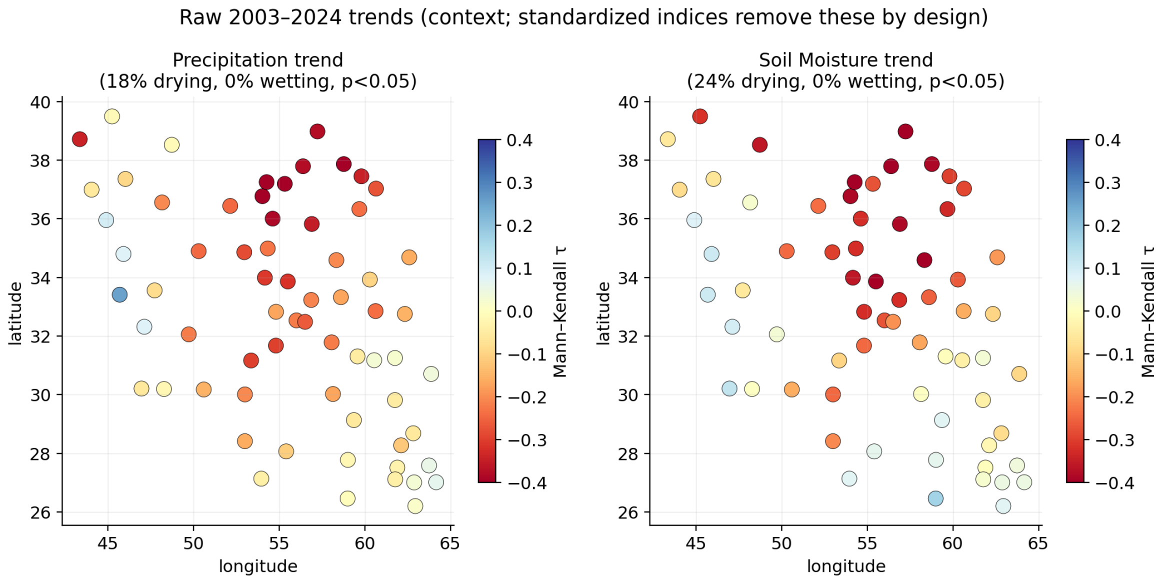

Figure 4. Mann–Kendall trend (τ) in raw annual precipitation (left) and surface soil moisture (right), 2003–2024. Red indicates drying. No basin shows a significant increase; drying concentrates in the northwest.

3.3 Fast surface response and its spatial pattern

The surface soil-moisture response is rapid and spatially organized, with the shortest timescales in the driest interior basins and longer timescales in wetter, more vegetated basins (Figure 5). The median accumulation timescale is two months (interquartile range 1–2 months); 38% of basins respond at a one-month accumulation and 94% within two months, while only 2% require four months or more (Figure 6). The best-fitting additional lag is zero months in every basin. This uniform result is expected rather than surprising: at monthly resolution a shallow (0–7 cm) layer that responds within days reaches its strongest correlation at the concurrent-month accumulation, so sub-monthly delays cannot be resolved, and the timescale should be read as the precipitation-integration window of the surface store rather than a literal travel time. The exact per-basin window is, moreover, only weakly identified: the |r| difference between the best and second-best accumulation window is small (median 0.02), and the top two windows lie within 0.05 of one another in 91% of basins, so the national two-month median is robust while individual one- versus two-month assignments are often within estimation noise.

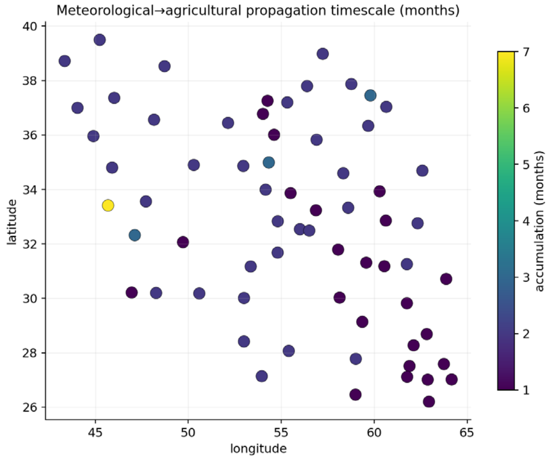

Figure 5. Meteorological-to-agricultural propagation timescale (accumulation period, months) for each sub-basin, plotted at basin centroids.

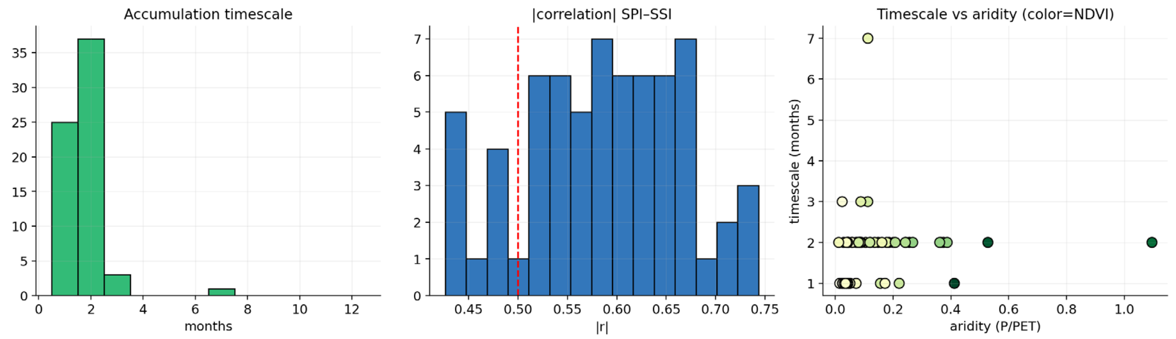

Figure 6. (left) Distribution of propagation timescales; (centre) distribution of |SPI–SSI correlation|, dashed line at |r| = 0.5; (right) timescale versus the P/PET humidity index, coloured by NDVI.

The estimated timescale is well constrained rather than an arbitrary maximum: the median coupling strength peaks at a one-to-two-month accumulation window and declines monotonically thereafter, with the across-basin range of |r| spanning roughly 0.24 between the best and worst windows (Figure 7). This confirms that the short propagation timescale is a genuine feature of the data.

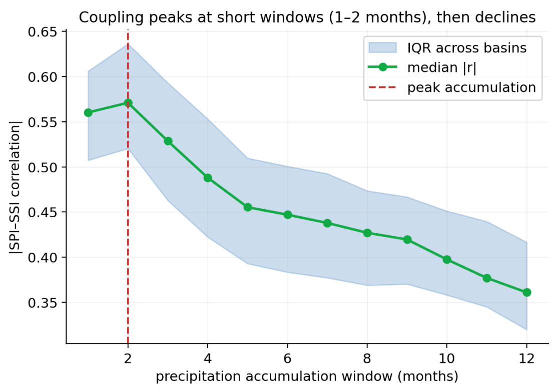

Figure 7. Median |SPI–SSI| correlation as a function of precipitation accumulation window, with the inter-quartile range across basins.

3.4 Descriptive associations with basin attributes

In-sample, the propagation timescale increases with cropland fraction (ρ = 0.49), NDVI (ρ = 0.49), the P/PET humidity index (ρ = 0.44; higher values denote more humid basins), and snow fraction (ρ = 0.42), and more weakly with population density (ρ = 0.36); elevation and the endorheic flag show no significant association (Table 3; Figure 8). Equivalently, the surface response is fastest in the driest, most water-limited basins and slower in wetter, more vegetated ones—physically intuitive, since thin soils and high evaporative demand transmit precipitation deficits rapidly (Huang et al., 2017). We note that Iran lies almost entirely within the water-limited regime (52 of 66 basins have P/PET < 0.2 and only two exceed 0.5), so this is a gradient within arid conditions rather than a transition between water- and energy-limited regimes. These associations are descriptive and are not corrected for multiple comparisons; they are not the inferential core of the study, which is the out-of-sample test that follows. As shown next, the pattern does not translate into out-of-sample predictability.

Table 3. Descriptive (in-sample) Spearman associations between propagation timescale and static basin attributes.

Attribute	Spearman ρ	p-value
Cropland fraction	+0.49	<0.001
NDVI	+0.49	<0.001
P/PET humidity index	+0.44	<0.001
Snow fraction	+0.42	<0.001
Population density	+0.36	0.003
Slope	+0.23	0.062
Permanent-water fraction	+0.21	0.092
Endorheic flag	−0.12	0.343
Elevation	+0.04	0.760

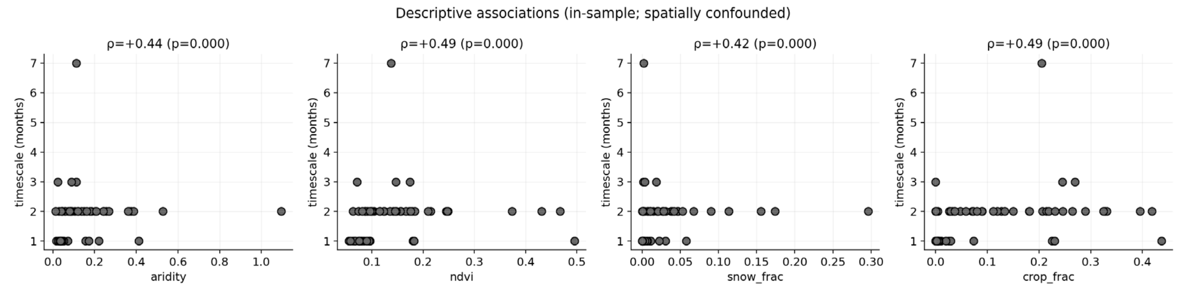

Figure 8. Descriptive (in-sample) associations between propagation timescale and the P/PET humidity index, NDVI, snow fraction, and cropland fraction. Spearman ρ and p-values shown; relationships are spatially confounded (see Section 3.6).

3.5 Limited predictability: random versus spatial cross-validation

Under random k-fold cross-validation the gradient-boosted model explains almost no variance ($R^2 = 0.08$) and the linear model none ($R^2 = -0.01$). Under spatial-block cross-validation, where whole regions are held out, skill is negative for both models (gradient-boosted $R^2 = -0.18$; linear $R^2 = -0.64$; Table 4; Figure 9); a negative R^2 indicates predictions worse than simply using the mean. The strong in-sample associations of Section 3.4 are therefore largely an expression of spatial autocorrelation—neighbouring basins resemble one another in both attributes and timescale—so a model can appear skilful in random cross-validation while failing entirely when whole regions are held out. We interpret this as evidence that propagation speed is set by transient hydroclimatic conditions rather than by fixed basin properties; the formal power and robustness checks supporting this interpretation are reported in Section 3.7.

Table 4. Cross-validated predictive skill (R^2) for the propagation timescale from static attributes. Negative R^2 denotes prediction worse than using the sample mean.

Model	Random-CV R^2	Spatial-block R^2
Gradient-boosted trees	+0.08	−0.18
Linear regression	−0.01	−0.64

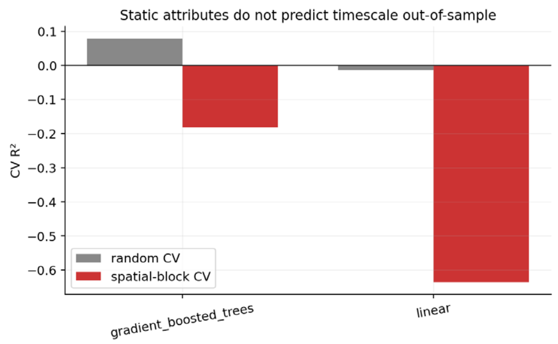

Figure 9. Cross-validated R^2 for predicting propagation timescale from static basin attributes, under random k-fold (grey) and spatial-block (red) cross-validation. Negative spatial-block skill indicates predictions worse than the mean.

3.6 Spatial structure and the limits of attribute-based prediction

The propagation timescale shows modest but statistically significant spatial autocorrelation (Moran's I = 0.17, $p = 0.002$): nearby basins tend to share somewhat similar response times, confirming that the in-sample associations of Section 3.4 are spatially confounded. Under identical spatial-block folds, a location-only inverse-distance neighbour model attains $R^2 = -0.02$ (median across partitions ≈ 0.00), materially better than the attribute model's $R^2 = -0.18$ but still no better than the no-skill mean (Table 5; Figure 10).

The correct reading is therefore not that location 'beats' attributes in any predictive sense—neither model has positive out-of-sample skill—but that the attribute model performs worse than the mean while location performs at the mean level. Location here is best understood as a proxy for spatially structured but unmeasured influences (transient climate states, sub-grid soil and groundwater properties) rather than as a causal control. The negative attribute skill is not an artefact of one blocking choice: it persists across four to ten spatial blocks (R^2 between −0.19 and −0.11). Collectively, the static descriptors examined here carry no generalizable, location-independent information about propagation speed that we can detect with the available sample. Spatial blocks were defined by k-means on basin centroids so that whole, geographically contiguous regions are held out; as shown above, the negative skill is insensitive to the number of blocks, so it does not depend on this particular partitioning.

Table 5. Spatial diagnostics: autocorrelation of the timescale, attribute vs location-only skill, and robustness to blocking.

Diagnostic	Value
Moran's I (timescale)	+0.17 ($p = 0.002$)
Spatial-CV R^2 — attribute model	−0.18
Spatial-CV R^2 — location-only neighbour model	−0.02
Attribute spatial-CV R^2, k = 4 / 6 / 8 / 10 blocks	−0.19 / −0.18 / −0.11 / −0.13
Index robustness (empirical-Gaussian): timescale, |r|	2 months, 0.60

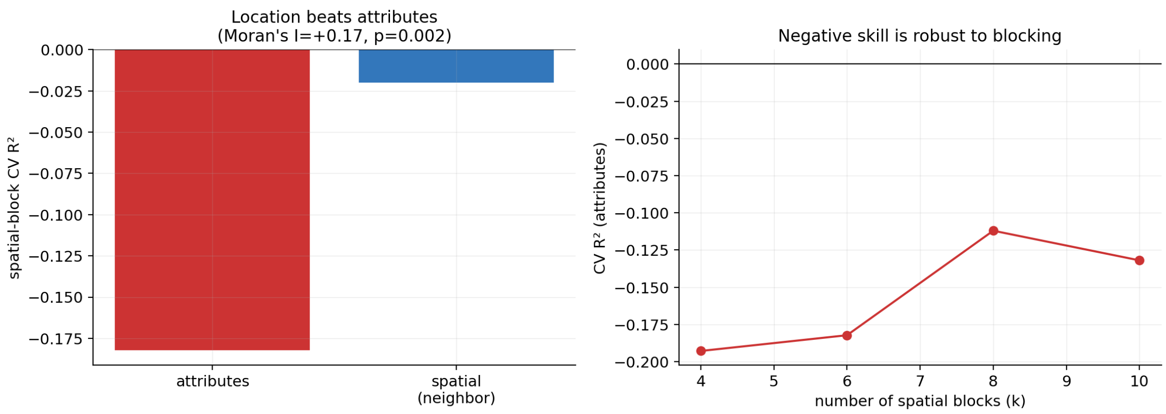

Figure 10. (left) Spatial-block cross-validated R^2 for predicting propagation timescale from static attributes versus from a location-only neighbour model; the timescale is significantly spatially autocorrelated. (right) The negative attribute skill is robust to the number of spatial blocks.

The findings are also insensitive to the index definition. Repeating the analysis with a non-parametric empirical-Gaussian standardization yields the same median timescale (two months) and a near-identical median coupling strength ($|r|$ = 0.60 versus 0.58), although individual basin timescales agree only moderately between the two methods (ρ = 0.44), a sensitivity we return to in the limitations.

3.7 Statistical power, robustness, and temporal stability

Three results show that the negative predictability is a genuine absence of transferable signal rather than an artefact of small sample size, and one shows that per-basin timescales are not temporally fixed (Table 6; Figure 11). First, across 50 independent spatial-block partitions the attribute model's R^2 is tightly and entirely negative (median −0.18, 95% interval [−0.19, −0.13]), whereas the location-only model straddles zero (median ≈ 0.00, [−0.05, +0.02]); the combined location-plus-attribute model reaches only +0.01, no better than location alone. Second, dimensionality reduction does not rescue attribute skill: spatial-CV R^2 from the first two to four principal components remains at or below zero (−0.01 to −0.08).

Third, the positive control succeeds—when the target is constructed to depend genuinely on an attribute, the identical spatial-CV procedure recovers R^2 = +0.19, +0.29, and +0.52 at weak, moderate, and strong signal levels (Figure 11, left). The procedure can therefore detect a real attribute signal at n = 66; its failure on the observed timescale reflects the absence of such a signal, not insufficient data.

Finally, the national median timescale is unchanged between halves (two months in both 2003–2013 and 2014–2024), but per-basin timescales do not reproduce between the two periods (Spearman ρ = −0.14, p = 0.26, i.e. not significant). We read this cautiously: a non-significant correlation provides no evidence that the per-basin timescale is a stable, transferable trait, but it does not by itself prove non-stationarity, because half-length records and the weak per-basin identifiability noted in Section 3.3 both inflate the disagreement. Either interpretation points the same way—per-basin timescales are not reliably fixed quantities—which is consistent with the failure of static-attribute prediction.

Table 6. Power, robustness, and temporal-stability checks.

Check	Result
Attribute spatial-CV R^2 (median [95% CI], 50 partitions)	−0.18 [−0.19, −0.13]
Location-only spatial-CV R^2 (median [95% CI])	0.00 [−0.05, +0.02]
Combined location + attributes R^2	+0.01
PCA-reduced R^2 (2 / 3 / 4 components)	−0.01 / −0.06 / −0.08
Positive control R^2 (weak / moderate / strong signal)	+0.19 / +0.29 / +0.52
Temporal median timescale (2003–13 / 2014–24)	2 / 2 months
Temporal per-basin rank correlation between halves	$\rho = -0.14$ ($p = 0.26$)

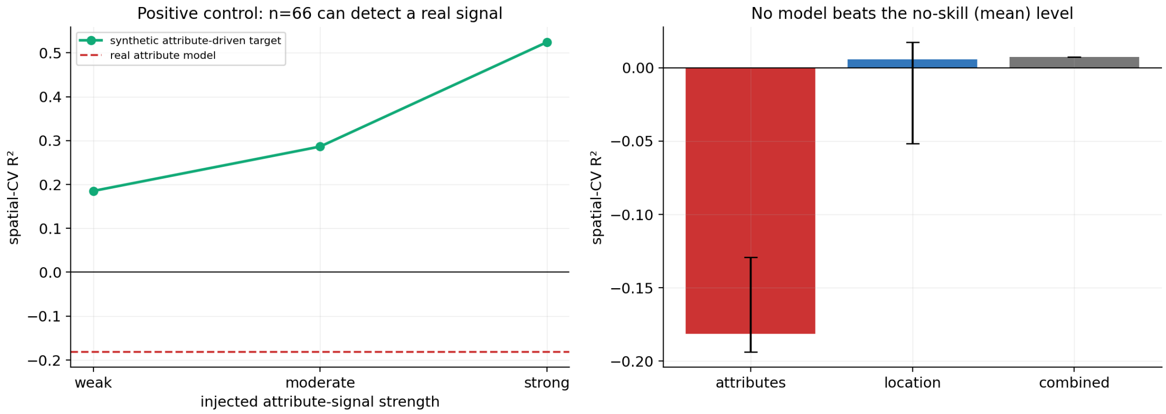

Figure 11. (left) Positive control: the identical spatial-CV procedure recovers increasing R^2 for synthetic targets built to depend on an attribute (green), well above the real attribute model (red dashed). (right) Spatial-CV R^2 with 95% intervals for the attribute, location-only, and combined models; none exceeds the no-skill (mean) level.

4 Discussion

Our findings confirm the well-established meteorological-to-agricultural drought cascade (Van Loon, 2015) and provide, to our knowledge, the first basin-resolved characterization of its timescale across all of Iran, complementing prior regional evidence from other arid and semi-arid settings with a national-scale assessment for a large, data-scarce country. Short meteorological-to-agricultural response times of roughly one to three months for surface soil moisture are consistent with prior multi-region findings; our median two-month timescale falls squarely in this range. Zhu et al. (2021), in a global analysis of meteorological-to-soil-moisture propagation probability, reported median response times of one to three months across water-limited regions, with faster propagation in more arid basins—a pattern we replicate nationally for Iran. Regional studies in semi-arid India (Choudhari et al., 2025) and northwest China (Li and Cai, 2024) report similar short lags (one to four months), reinforcing that the two-month median we observe is characteristic of water-limited environments rather than an Iran-specific anomaly. Satellite-derived SSI has been shown to capture drought dynamics effectively in Mediterranean settings (Escorihuela et al., 2025), further validating the cross-regional applicability of the SPI–SSI framework. Throughout, we stress that these timescales represent the near-surface (0–7 cm) soil-moisture response rather than the full root-zone response that defines agricultural drought in the agronomic sense; the surface store is a rapid, observable proxy for the onset of agricultural water stress rather than a measure of root-zone depletion.

Two findings stand out. First, the meteorological-to-agricultural link in Iran is strong and fast almost everywhere: a one-to-two-month precipitation accumulation explains surface soil-moisture anomalies in the large majority of basins, with significant correlations after a conservative autocorrelation correction. For drought early warning this implies that surface soil-moisture drought can be anticipated from short-window precipitation monitoring across most of the country, including data-scarce interior basins. This aligns with machine-learning drought studies that consistently rank SPI among the top predictors of hydrological drought across diverse regions (Yao, 2024).

Second, and more cautionary, the spatial variation of this timescale is not predictable from the standard static basin descriptors examined here once spatial cross-validation is applied; neither an attribute model, a location-only model, nor their combination exceeds the no-skill level, and the attribute model is in fact worse than the mean. Combined with the significant Moran's I and the successful positive control, this is direct evidence that the spurious in-sample associations are an expression of spatial autocorrelation rather than causal control: neighbouring basins resemble one another in both attributes and timescale, so a model can appear skilful in random cross-validation while failing entirely when whole regions are held out. Because the same procedure recovers a clear signal from a synthetic attribute-driven target at the same sample size, the null result cannot be attributed merely to the modest number of basins. We argue that reporting the spatial-block result, the neighbour benchmark, the positive control, and the blocking-robustness check transparently is more useful than an over-fitted attribution that would not survive replication in a new region.

Substantively, the result implies that the surface propagation timescale is set by transient hydroclimatic conditions—the sequencing, intensity, and phase of wet and dry spells—and by broad regional structure, rather than by time-invariant basin properties. This echoes the finding of Jehanzaib et al. (2023) that hydrological non-stationarity and environmental change can modify propagation pathways, and the observation of Liu et al. (2024) that human intervention can weaken propagation signals; the value of spatially explicit validation in water-limited regions is likewise demonstrated by Bouaicha et al. (2024). The exclusion of the hydrological (storage) stage is a scope decision driven by data physics: at HydroBASINS level 5, basins are smaller than the GRACE mascon footprint, so a per-basin storage analysis would conflate neighbouring basins. A defensible hydrological-stage analysis would require coarser analysis units (for example, HydroBASINS level 3–4) whose areas match or exceed the GRACE resolution.

These statistical results invite a process-based interpretation. A precipitation anomaly reaches the surface soil-moisture store through infiltration, near-surface storage and redistribution, vegetation uptake, and evapotranspiration, and the speed of that transmission depends on the antecedent wetness state, the intensity and phasing of individual storms, and the evaporative demand prevailing at the time—quantities that vary from event to event and season to season. Time-invariant descriptors such as aridity, vegetation cover, or cropland fraction summarize the long-term setting but cannot encode this dynamic state. The out-of-sample failure of those descriptors is therefore physically consistent with propagation speed being governed mainly by transient hydroclimatic conditions, with the static attributes acting only as weak, spatially organized correlates rather than as controlling variables.

That per-basin timescales do not reproduce between the two halves of the record offers a complementary reading of the predictive null: if a basin's timescale cannot be pinned down consistently over time—whether because it genuinely varies or because it is only weakly identified at half-record length—then it is unlikely to be a fixed function of that basin's time-invariant attributes. This non-reproducibility is itself a noteworthy result: it cautions against the common assumption that each basin possesses a single characteristic propagation timescale, and argues for time-resolved rather than static characterizations and for updating operational propagation estimates as conditions evolve rather than fixing them once per basin.

4.1 Implications and applications

The one-to-two-month propagation timescale has direct implications for drought early warning. Surface soil moisture directly influences germination, crop establishment, and early-season water stress, so these response times set the warning horizon available for agricultural drought preparedness. Because precipitation deficits are transmitted to the near-surface soil layer within one to two months across most Iranian basins, routine monitoring of short-window precipitation anomalies (for example, one- to two-month SPI) can provide useful lead time for anticipating surface soil-moisture drought—the onset of agricultural water stress—and thereby support irrigation scheduling, planting decisions, and seasonal risk assessment, including across the data-scarce interior where in-situ observations are sparse. By the same logic, the basins that respond at a one-month accumulation are the most exposed to short-lived precipitation deficits coinciding with sensitive crop-growth stages, and may warrant closer operational monitoring. Because the response measured here is that of the surface layer, these implications concern the early, near-surface phase of agricultural drought rather than deep root-zone depletion.

Although the analysis is hydroclimatic, the surface soil-moisture response timescales are also relevant to geotechnical processes driven by near-surface moisture fluctuations, including shrink–swell behaviour, desiccation cracking, and seasonal stiffness variations in expansive soils. The one-to-two-month response window provides a first-order estimate of the timescale over which meteorological anomalies may influence near-surface engineering soil conditions at regional scales. We note this only as a potential application, since the present study includes no suction, swelling, settlement, or other geotechnical measurements.

4.2 Limitations

The analysis rests on 66 basins, which is ample for the correlation-based characterization but modest for predictive modelling; we have therefore reported R^2 uncertainty intervals, a positive control, a combined model, and a dimensionality-reduced model, which together make the spatial-confounding conclusion secure rather than a possible power artefact. Soil moisture is taken from ERA5-Land, a model-based reanalysis rather than a direct observation, and represents only the surface 0–7 cm layer; we accordingly characterize the surface soil-moisture response and interpret the SSI as a surface agricultural-drought proxy. Deeper root-zone dynamics (roughly 30–100 cm) relevant to many crops are not resolved, and the reported timescales should be read as surface response times, which may be shorter than full root-zone response times. While aggregate conclusions are insensitive to the standardization method, individual basin timescales agree only moderately between the z-score and empirical-Gaussian definitions (ρ = 0.44) and between record halves (ρ = −0.14), so per-basin timescale values should be treated as indicative rather than exact (Section 3.3).

The attribute set is also deliberately limited to remotely sensed and topographic descriptors; potentially important controls were unavailable, including soil texture and depth, saturated hydraulic conductivity, groundwater depth and aquifer type, and sub-monthly rainfall characteristics such as intermittency and storm intensity. The null result therefore applies to the attributes examined here and does not exclude predictability from richer soil-hydrological or sub-monthly climate predictors; repeating the analysis with root-zone products such as SMAP Level-4 or GLEAM soil moisture would test whether the surface response times reported here extend to the agronomically relevant root zone. Finally, the timescale is estimated over the full record; because Iran's climate is drying (Section 3.2) and propagation dynamics may shift as aridity intensifies (Jehanzaib et al., 2023; Cammalleri et al., 2022), reviews of drought assessment under changing climate normals caution that fixed-reference standardization may obscure evolving relationships (Lisonbee et al., 2024). A continuous rolling-window analysis is a natural extension of the split-half stability test reported here.

5 Conclusion

Across Iran's 66 level-5 sub-basins, meteorological drought propagates into the surface soil-moisture store strongly and rapidly: significant SPI–SSI coupling in every basin after autocorrelation correction, a median correlation of 0.58, and a median surface response timescale of two months. Although the timescale varies coherently with aridity and vegetation in-sample, it is significantly spatially autocorrelated and cannot be predicted from the static basin attributes examined here out-of-sample; no model, including a location-only benchmark, exceeds the no-skill level, while a positive control confirms that a genuine attribute signal would be detectable at this sample size. Per-basin timescales, moreover, do not reproduce between sub-periods, giving no evidence that they are fixed, transferable quantities. Propagation speed in this region is therefore an emergent property of transient hydroclimate and spatial structure, not a fixed basin trait. The study delivers a satellite-and-reanalysis national baseline for surface agricultural-drought response timing, a transparent validation framework—spatial cross-validation, a spatial null model, a positive control, and temporal-stability testing—that exposes spatial-autocorrelation pitfalls, and a fully reproducible workflow for extending the approach to coarser analysis units and to the hydrological stage.

Acknowledgments

The authors acknowledge the providers of the publicly available datasets used in this study, including the Climate Hazards Center for the CHIRPS precipitation dataset, the European Centre for Medium-Range Weather Forecasts (ECMWF) for the ERA5-Land reanalysis, and the relevant providers of the satellite, land-cover, elevation, surface-water, and population datasets used in the analysis. The authors also acknowledge the Google Earth Engine platform for providing access to and processing of these openly available geospatial datasets. No external funding was received for this research, and the study was conducted without financial support or external research assistance.

Data and code availability

All input variables were generated in Google Earth Engine and exported as basin-level monthly and static tables. The complete analysis is implemented in a single reproducible Python script that regenerates every table and figure in this manuscript from the input CSV files and produces a complete machine-readable results file. The input data, analysis code, and results files are available from the corresponding author upon reasonable request.

References

1. Bouaicha, R., Lahlou, A. A., Strobl, E., et al. (2024). Nested cross-validation for HBV conceptual rainfall–runoff model spatial stability analysis in a semi-arid context. *Remote Sensing*, 16(20), 3756.
2. Bretherton, C. S., Widmann, M., Dymnikov, V. P., Wallace, J. M., & Bladé, I. (1999). The effective number of spatial degrees of freedom of a time-varying field. *Journal of Climate*, 12(7), 1990–2009.
3. Cammalleri, C., Spinoni, J., Barbosa, P., Toreti, A., & Vogt, J. V. (2022). The effects of non-stationarity on SPI for operational drought monitoring in Europe. *International Journal of Climatology*, 42(6), 3418–3430.
4. Choudhari, N., Jacob, B. G., Elshorbany, Y., & Collins, J. (2025). Unveiling asymptotic behavior in precipitation time series. *Hydrology*, 12(10), 254.
5. Escorihuela, M. J., Pablos, M., Quintana-Seguí, P., et al. (2025). Remote sensing standardized soil moisture index for drought monitoring: a case study in the Ebro Basin. *Remote Sensing*, 17(23), 3916.
6. Farahmand, A., AghaKouchak, A., & Ek, M. B. (2021). Standardized soil moisture index for drought monitoring based on SMAP observations and 36 years of NLDAS data: a case study in the southeast United States. *Remote Sensing of Environment*, 256, 112319.
7. Funk, C., Peterson, P., Landsfeld, M., et al. (2015). The climate hazards infrared precipitation with stations—a new environmental record for monitoring extremes. *Scientific Data*, 2, 150066.
8. Huang, S., Li, P., Huang, Q., Leng, G., Hou, B., & Wei, G. (2017). The propagation from meteorological to hydrological drought and its potential influencing factors. *Journal of Hydrology*, 547, 184–195.
9. Jehanzaib, M., Ali, S., Kim, M. J., et al. (2023). Modeling hydrological non-stationarity to analyze environmental impacts on drought propagation. *Atmospheric Research*, 286, 106699.
10. Jiang, T., Su, X., Zhang, G., et al. (2023). Estimating propagation probability from meteorological to ecological droughts using a hybrid machine-learning copula method. *Hydrology and Earth System Sciences*, 27(2), 559–576.
11. Khoorani, A., Balaghi, S., & Mohammadi, F. (2024). Projecting drought trends and hot spots across Iran. *Natural Hazards*, 120, 9489–9502.
12. Kumar, V., & Chu, H.-J. (2024). Spatiotemporal consistency and inconsistency of meteorological and agricultural drought identification: a case study of India. *Remote Sensing Applications: Society and Environment*, 33, 101134.
13. Lehner, B., & Grill, G. (2013). Global river hydrography and network routing: baseline data and new approaches to study the world's large river systems. *Hydrological Processes*, 27(15), 2171–2186.
14. Li, L., & Cai, H. (2024). A comparative study of various drought indices at different timescales and over different record lengths in the arid area of northwest China. *Environmental Science and Pollution Research*, 31, 25096–25113.
15. Lisonbee, J., Nielsen-Gammon, J., Trewin, B., Follingstad, G., & Parker, B. (2024). Drought assessment in a changing climate: a review of climate normals for drought indices. *Journal of Applied Service Climatology*, 2024, 1–17.
16. Liu, X., Peng, J., Liu, Y., et al. (2024). The Three Gorges Dam has weakened the drought propagation process in the Yangtze River Basin. *Journal of Hydrology*, 632, 130875.
17. McKee, T. B., Doesken, N. J., & Kleist, J. (1993). The relationship of drought frequency and duration to time scales. *Proceedings of the 8th Conference on Applied Climatology*, 179–184.

18. Muñoz-Sabater, J., Dutra, E., Agustí-Panareda, A., et al. (2021). ERA5-Land: a state-of-the-art global reanalysis dataset for land applications. *Earth System Science Data*, 13(9), 4349–4383.

19. Ploton, P., Mortier, F., Réjou-Méchain, M., et al. (2020). Spatial validation reveals poor predictive performance of large-scale ecological mapping models. *Nature Communications*, 11, 4540.

20. Roberts, D. R., Bahn, V., Ciuti, S., et al. (2017). Cross-validation strategies for data with temporal, spatial, hierarchical, or phylogenetic structure. *Ecography*, 40(8), 913–929.

21. Van Loon, A. F. (2015). Hydrological drought explained. *WIREs Water*, 2(4), 359–392.

22. Vicente-Serrano, S. M., Beguería, S., & López-Moreno, J. I. (2010). A multiscalar drought index sensitive to global warming: the standardized precipitation evapotranspiration index. *Journal of Climate*, 23(7), 1696–1718.

23. World Weather Attribution. (2024). Climate change: intensity of ongoing drought in Syria, Iraq and Iran 'not rare anymore'. https://www.worldweatherattribution.org

24. Xu, L., Zhang, X., Yu, H., Chen, Z., Du, W., & Chen, N. (2024). Incorporating spatial autocorrelation into deformable ConvLSTM for hourly precipitation forecasting. *Computers & Geosciences*, 184, 105536.

25. Yao, Y. (2024). Hydrological drought prediction and its influencing features in the Huaihe River Basin. *Natural Hazards and Earth System Sciences*, 25, 4299–4320.

26. Zhu, Y., Liu, Y., Wang, W., et al. (2021). A global perspective on the probability of propagation of drought: from meteorological to soil moisture. *Journal of Hydrology*, 603, 126907.