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.

Large-scale causal drivers of European winter precipitation

Marc Peradalta-Negre

marc.peradalta@upc.edu

Universitat Politecnica de Catalunya https://orcid.org/0009-0004-3976-3352

Jonathan Ortved Melcher

University of Copenhagen Niels Bohr Institute: Kobenhavns Universitet Niels Bohr Instituttet

Andrea Vang

University of Copenhagen Niels Bohr Institute: Kobenhavns Universitet Niels Bohr Instituttet

Jens Hesselbjerg Christensen

University of Copenhagen Niels Bohr Institute: Kobenhavns Universitet Niels Bohr Instituttet

Ramiro I. Saurral

Barcelona Supercomputing Center: Centro Nacional de Supercomputacion

Research Article

Keywords: Causal inference, Winter precipitation, Teleconnections, North Atlantic Oscillation

Posted Date: November 11th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-11469447/v1

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

Large-scale causal drivers of European winter precipitation

Marc Peradalta–Negre[1,2,3*], Jonathan Ortved Melcher[2], Andrea Vang[2,4], Jens Hesselbjerg Christensen[2], Ramiro I. Saurral[3]

[1*]Department of Mathematics, Universitat Politècnica de Catalunya (UPC), Barcelona, Spain.
[2]Physics of Ice, Climate and Earth, Niels Bohr Institute, University of Copenhagen, Copenhagen, Denmark.
[3]Barcelona Supercomputing Center (BSC), Barcelona, Spain.
[4]National Center for Climate Research, Danish Meteorological Institute, Copenhagen, Denmark.

*Corresponding author(s). E-mail(s): marc.peradalta@upc.edu, ORCID: 0009-0004-3976-3352;
Contributing authors: jonathan.melcher@nbi.ku.dk, ORCID: 0009-0006-9473-1200; andrea.vang@nbi.ku.dk, ORCID: 0009-0004-2142-5654; hesselbjerg@nbi.ku.dk, ORCID: 0000-0002-9908-8203; ramiro.saurral@bsc.es, ORCID: 0000-0002-8600-199X;

Abstract

Large-scale atmospheric dynamics drive a substantial share of winter precipitation in the Euro-Atlantic sector. Still, the associated mechanisms remain only partially understood despite their relevance for seasonal forecasting and climate predictions and projections. While averaged modes like the North Atlantic Oscillation (NAO) prove useful in regressing precipitation, in this study we aim to add a layer of physical interpretability and disentangle individual mechanisms that affect precipitation. Using a causal discovery framework applied independently at every grid point across the Euro-Atlantic domain, we identify direct large-scale drivers of winter precipitation locally, with dependencies found to be predominantly linear and contemporaneous at monthly timescales. Causal links are uncovered, including a direct connection between Barents–Kara sea ice and Iberian Peninsula precipitation, together with evidence that the stratospheric polar vortex causally associates with precipitation across northern and southern

Europe without requiring other large-scale intermediaries. Unsupervised clustering of the identified causal structure organises local precipitation into coherent regions resolving most of the Euro-Atlantic domain. Using the identified causal drivers, a multiple linear regression framework on precipitation achieves significant explanatory power over a substantially larger fraction of the domain than the NAO, and regionally outperforms it in explained variance away from its centres of action, showing that a physically grounded understanding of the mechanisms driving precipitation also yields a better description of it.

Keywords: Causal inference, Winter precipitation, Teleconnections, North Atlantic Oscillation

1 Introduction

Precipitation is a key variable of the climate system that controls freshwater availability, agricultural productivity, and the frequency and intensity of floods and droughts (Caretta et al. 2022). It is strongly modulated by the atmospheric circulation, which determines regions and periods in which precipitation is fostered or hindered through the onset of blocking anticyclones (Wang et al. 2026), the pathways of extratropical cyclones and the storm tracks (Tamarin-Brodsky et al. 2026), and the location and evolution of atmospheric rivers (Sneddon et al. 2026), among others. Even though the characteristics of these weather features are well known for the past and present climate, their response to global warming is a problem not well constrained, particularly for the midlatitudes (Shepherd 2014).

This is the case of Europe, where dynamically-driven (i.e. large-scale) precipitation peaks during the cold season (Berg et al. 2009). The main mode of winter atmospheric circulation over that region is the North Atlantic Oscillation (NAO, Rogers 1984), which correlates significantly with precipitation across the region (Hurrell 1995). This mode, identified in surface pressure and geopotential height fields at low, mid and upper tropospheric levels, can be broadly described as a pressure dipole between the British Isles and the Azores region. A direct consequence of the location, polarity and strength of this dipole is the presence of rainfall anomalies of opposite sign over northern Europe (British Isles, Scandinavia) and southern Europe (Portugal, Spain, southern France). However, the NAO can also be treated as the summarised effect from many different mechanisms acting contemporaneously on the North Atlantic region and Europe, and following this approach might offer extended interpretability into the individual drivers, and future changes thereof. For example, Arctic-to-midlatitude teleconnections have received particular attention in recent literature, motivated by the rapid decline of Arctic sea ice and its potential to alter mid-latitude circulation and precipitation (Cohen et al. 2014b; Saurral et al. 2025), highlighting the importance of understanding these pathways to assess the full impact of climate change on precipitation.

Identifying these teleconnections from observational data is, however, a non-trivial task. The climate system is characterised by high-dimensional interactions and non-linearities (Palmer 1999; Hannachi et al. 2017) that make it difficult to isolate cause and effect from statistical associations alone (Ebert-Uphoff and Deng 2012; Kretschmer et al. 2021). Traditional correlation-based approaches are prone to identifying spurious links, and the NAO is itself an example: it correlates well with precipitation not because it causes it in a physically direct sense, but because it aggregates the effects of many drivers into a single index (Woollings et al. 2010). In this regard, causal inference offers a statistically grounded framework to go beyond correlation, distinguishing direct causal influences from indirect or confounded ones (Kretschmer et al. 2016; Runge et al. 2019a).

In this study we apply a causal discovery algorithm over a set of candidate drivers at every grid point across the Euro-Atlantic domain to (i) identify the direct large-scale drivers of local winter precipitation, (ii) test whether grid points sharing the same causal structure organise into spatially coherent regions, and (iii) quantify the regression skill of these drivers relative to the NAO. Together, these three elements provide a causally grounded characterisation of Euro-Atlantic winter precipitation that is both more physically interpretable and, as we show, with higher explanatory power than the NAO alone.

2 Data

2.1 Characteristics of the data

The study period spans 1979–2025. We restrict the analysis to the boreal winter season, defined as the December–January–February (DJF) trimester, when large-scale atmospheric circulation dominates precipitation variability in the Euro-Atlantic sector (Berg et al. 2009). All variables are aggregated into monthly means and expressed as anomalies relative to their monthly climatology; this averaging filters out high-frequency synoptic variability and focuses the analysis on the sub-seasonal to seasonal timescales at which teleconnections operate (Branstator and Teng 2017). To account for lagged effects, and since precipitation is evaluated over DJF, the candidate drivers are evaluated at leading months 0 to 3: a lag of $\tau = 0$ corresponds to DJF, $\tau = 1$ to NDJ, $\tau = 2$ to OND, and $\tau = 3$ to SON. This lead-lag structure allows the causal inference algorithm to detect not only contemporaneous associations but also precursor signals that anticipate winter precipitation.

The spatial domain covers the Euro-Atlantic sector (80°W–40°E, 30°N–85°N; Figure 1), encompassing the two centres of action of the NAO over the North Atlantic and its principal region of influence over the European continent. Most variables are derived from the ERA5 reanalysis (Hersbach et al. 2020), which combines observations with a numerical weather prediction model through data assimilation to provide a physically consistent, spatially and temporally complete estimate of the atmospheric state. The target variable, Total Precipitation (TP), retains its native 0.25° × 0.25° horizontal resolution, so that the causal discovery algorithm is applied independently at every grid point; the candidate drivers are instead scalar indices, each representative of a specific large-scale dynamical feature. The NAO (Barnston and

Livezey 1987) and El Niño–Southern Oscillation (ENSO, Barnston et al. 1997) indices are the two exceptions to the ERA5-based approach, taken instead from pre-computed, community-standard records available from the NOAA archive (Trenberth and Staff 2026; Hurrel et al. 2026).

2.2 Candidate drivers

Large-scale drivers influence Euro-Atlantic precipitation through atmospheric teleconnections, primarily mediated by the propagation of planetary Rossby waves (Hoskins and Karoly 1981), which carry the signals of the large-scale mechanisms potentially far from the source and onto precipitation. A *driver* (also referred to as a cause or link), defined here as a large-scale variable which influences precipitation, can be further classified as *direct* if it modulates precipitation without any other driver acting as an intermediary, or as *indirect* if it affects precipitation by modulating a direct driver. For example, the Polar Front (PF) might affect precipitation directly, while ENSO might be an indirect cause, by first modulating the PF, which in turn affects precipitation. This distinction is central to the causal discovery framework introduced in Section 3, which is designed to retain only direct links.

The *candidate* drivers – those used in our algorithm as potential direct drivers of precipitation – are selected following the broad classification of Hall et al. (2015), by choosing one or two representative indices for each of four driver categories. Figure 1 shows their corresponding regions of definition relative to the Euro-Atlantic domain; the SPV and PEHF are defined over all longitudes and are only partially displayed, while ENSO lies outside the map and is not shown.

Tropospheric circulation is represented by the Polar Front (PF) and the Ural Blocking (UB). The former is defined as the mean latitude of the 5520 m geopotential height contour at 500 hPa computed for each longitude in the band 50°W–40°E over the latitude range 30°N–85°N (Sanders 1988; Vang et al. 2026), while the latter is defined as the mean sea-level pressure over the Ural region (40°E–85°E, 45°N–70°N) (Cohen et al. 2014a). The PF tracks the storm track and is therefore a direct modulator of precipitation (Hall et al. 2015; Shaw et al. 2016; Priestley and Catto 2022). The UB, on its positive phase, deflects the storm track through its anticyclones, producing prolonged precipitation anomalies over Europe (Kautz et al. 2022).

Stratospheric dynamics is represented by the Stratospheric Polar Vortex (SPV) and the Poleward Eddy Heat Flux (PEHF). The SPV index is defined as the negative of the mean geopotential height anomaly averaged over all pressure levels between 10 and 100 hPa over the polar cap (65°N–90°N) (Kim et al. 2014; Cohen et al. 2014a). The PEHF is defined as $\overline{v'T'}$ at 100 hPa over the latitude band (45°N–75°N), where v is the meridional velocity, T is the temperature, primes denote deviations from the zonal mean, and the overbar denotes the zonal mean (Kim et al. 2014; Cohen et al. 2014a). Disruptions of the SPV propagate downward into the troposphere over subsequent weeks to months, driving circulation anomalies over Europe (Monnin et al. 2022; Hanna et al. 2024). The PEHF, proportional to the vertical component of the Eliassen–Palm flux, quantifies this vertical wave activity, acting as a mediator carrying information downwards (Hannachi et al. 2025).

Barents–Kara Seas Sea Ice Concentration (BKS-SIC), defined over the region 30°E–105°E, 70°N–80°N, is included as a cryospheric driver as it has been linked to Arctic Amplification (Cohen et al. 2014b) and to mid-latitude circulation changes (Kretschmer et al. 2016; McKenna et al. 2018; Smith et al. 2019; Hanna et al. 2024; Saurral et al. 2025; Ye et al. 2025; Saurral and Kucharski 2026).

Tropical forcing is represented by ENSO, quantified through the Niño 3.4 SST anomaly (5°S–5°N, 170°W–120°W). Its influence on Euro-Atlantic precipitation is well established but expected to be indirect, acting through intermediate processes rather than as a direct cause (Brönnimann 2007).

The NAO index is retained throughout as a reference statistical descriptor rather than a candidate driver: by comparing precipitation regressed onto the NAO with precipitation regressed onto the physically and causally motivated drivers above, we quantify how much of the explanatory power conventionally attributed to the NAO can instead be recovered through more interpretable mechanisms.

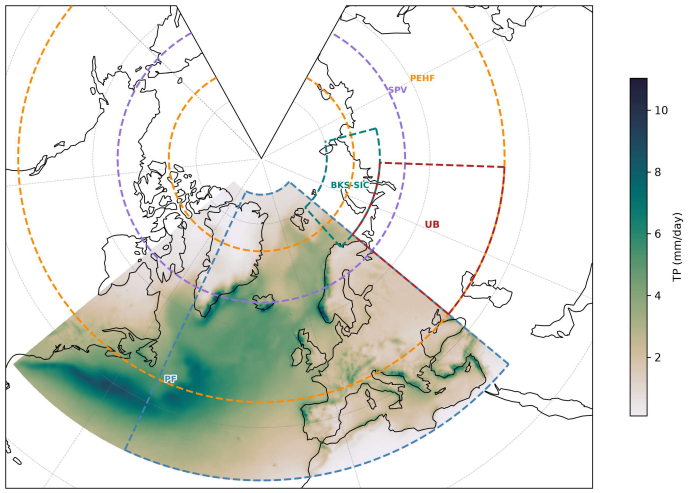


Fig. 1 Regions of definition of the candidate large-scale drivers (dashed lines). Shading shows the winter-mean (DJF) total precipitation (TP, mm day^{-1}) over the Euro-Atlantic domain. SPV and PEHF are defined over all longitudes, so for visual clarity only their latitude limits are shown, cut off at the edge of the map. ENSO (Niño 3.4) lies outside the map and is not shown.

3 Methods

3.1 Causal discovery framework

The key principle underlying causal discovery algorithms is that they remove spurious associations by conditioning on other potential drivers (Pearl 2009; Peters et al. 2017): a candidate link between two variables is retained only if the association between them persists after accounting for the influence of all other features that could otherwise explain it, distinguishing common-driver and indirect-driver configurations from genuine direct causal links (Runge et al. 2019a, 2023). Mathematically, a causal link can

be established via Conditional Independence (CI), with the formal bridge between the two explained in-depth in Runge et al. (2019b). The algorithm employed in this study is a variant of the PCMCI method introduced by Runge et al. (2019b) adapted from Kretschmer et al. (2016). It proceeds in two main stages: a Peter-Clark (PC) stage, which identifies a reduced set of potential drivers for each variable, and a Momentary Conditional Independence (MCI) stage, which performs the final CI tests conditioned on these potential drivers.

The principal methodological novelty of this paper relative to previous causal teleconnection studies is that the analysis is performed at every grid point of the Euro-Atlantic precipitation field, rather than on a small number of pre-defined regional indices. To do so, we designate TP as the target variable and run the algorithm independently at every grid point with the same set of candidate large-scale drivers introduced in Section 2, obtaining the local causal links directed towards precipitation.

Causal inference rests on a set of formal assumptions (Runge 2018), three of which are worth commenting here. Causal sufficiency cannot be guaranteed, but any unexplained variability can instead be read constructively as evidence of mechanisms lying outside our candidate driver set. Time ordering is enforced by construction through the use of lagged variables, except at $\tau = 0$ where contemporaneous links arise because our monthly resolution is coarser than some teleconnection timescales; we retain these and assign their directionality on physical grounds, since local precipitation cannot plausibly drive a large-scale atmospheric feature, so any such link reflects an influence on precipitation operating on submonthly scales. Causal stationarity is the least straightforward assumption to satisfy. Following Kretschmer et al. (2016), we tested the sensitivity of the algorithms to trends by repeating the analysis with BKS-SIC, the only variable exhibiting a significant trend, linearly detrended. The causal links obtained with and without detrending closely agree (not shown), so we retain the non-detrended climatological anomalies in the interest of physical simplicity and interpretability.

A key strength of the algorithm is its flexibility: the PC and MCI stages define its conditioning structure, but the CI tests at their core can be freely adapted to reflect different assumptions about the distributions and dependencies in the data. As mentioned above, a causal link between X and Y exists if, given the potential confounders $\mathbf{Z} = \{Z_1, ..., Z_m\}$, X and Y are dependent conditioned on $\mathbf{Z}$. This is a rather abstract condition, but CI can be translated into a computable quantity through a two-step procedure. The first step, *residualisation*, regresses out $\mathbf{Z}$ from each variable,

$$\begin{cases} X = f(\mathbf{Z}) + \varepsilon_X, \\ Y = g(\mathbf{Z}) + \varepsilon_Y, \end{cases} \tag{1}$$

so that the residuals ε_X and ε_Y retain the variability of X and Y that is not explained by $\mathbf{Z}$. The second step involves calculating the *marginal dependence of the residuals*, $\mathrm{Corr}(\varepsilon_X, \varepsilon_Y)$, where Corr is a correlation estimator. If $\mathrm{Corr}(\varepsilon_X, \varepsilon_Y) \neq 0$, we then say that a causal link exists between X and Y. In this study we employ a linear variant of the algorithm, namely *partial correlation* (Baba et al. 2004); f and g are chosen to be linear functions, so that residualisation translates into a multiple linear regression, followed by a Pearson correlation, ρ, on the residuals. Statistical significance is assessed

via a two-sided t-test with $n - 2 - m$ degrees of freedom (Fisher 1924), with $n \approx 140$ being the total number of samples and m the number of conditioning variables. The significance level is set to $\alpha = 0.05$, so that a partial correlation is deemed significant if its p-value satisfies $p \leq \alpha$. The other free parameter of the algorithm is the maximum lag $\tau_{\max} = 3$ months, reflecting the characteristic timescales of the teleconnections under study (Branstator and Teng 2017): any causal influence of a large-scale driver on precipitation should manifest within three months, otherwise it would be too noisy or mediated by intermediate mechanisms. While $\tau_{\max}$ is imposed on physical grounds, the significance level $\alpha = 0.05$ was chosen after testing on a range of values. Stricter levels discard a large fraction of links and leave many grid points with no identified driver, which limits the physical interpretation of the results, whereas more permissive levels retain a growing number of weak links that are more likely to be spurious.

The algorithm yields, at every grid point, a local set of causal links describing which large-scale drivers influence precipitation at that location and at which lag, emphasizing that a causal link is the combination of a large-scale driver at a specific lag. This means, for example, that PF might be a driver of TP at a specific grid point for lag $\tau = 2$, but might not for $\tau = 1$.

Once the causal links set has been identified at every grid point, we apply k-means clustering to group grid points that share the same large-scale causal drivers. Critically, clustering is performed on the causal driver structure itself – each grid point is represented as a vector encoding the strength of its causal link to each candidate variable-lag pair – rather than on the precipitation field or on geographical coordinates. Any spatial coherence among the resulting clusters therefore emerges purely from similarities in causal structure. We use $k = 5$ throughout the results presented below, chosen after inspecting the elbow of the inertia curve and the resulting clustered maps for $k \in \{4, ..., 9\}$ (not shown).

3.2 Regression framework

We quantify the explanatory power of the identified drivers using multiple linear regression at both the local and regional scale. At the local scale, for each grid point we regress precipitation onto its causal drivers identified by PCMCI and quantify their *explanatory power* through the R^2 of the regression, hereafter also referred to as the *regression skill*. The NAO benchmark is defined analogously, using the NAO index at all lags as regressors. Statistical significance is assessed as follows. For grid points with a single causal driver, a two-sided t-test is used, as in the linear CI test. For grid points with more than one driver, an F-test is used instead, which tests the significance of all regressors simultaneously.

At the regional scale, precipitation is averaged over all grid points within each cluster, weighting by grid-cell area, and the cluster-averaged time series is regressed onto the causal drivers identified at the corresponding cluster centroid, retaining only the variable-lag pairs with a link strength of $|\rho| \geq 0.1$. This threshold is used to reduce the dimensionality of the regression by discarding pairs with negligible link strength at the centroid, which would add regressors with little explanatory contribution given the limited sample size. The NAO index regressed onto the same cluster-averaged precipitation is used again as a benchmark.

4 Results

4.1 Drivers of precipitation

Figure 2 shows the causal link strength, or partial correlation, between precipitation and each of PF, SPV, UB, and BKS-SIC at $\tau = 0$ and $\tau = 1$ (full lag structure, including ENSO and PEHF, is provided in Appendix A, Figure A1). As expected, no significant direct link is found between precipitation and either ENSO or PEHF at any lag, consistent with both variables acting as indirect or mediating influences rather than direct causes. On average, we identify 2.69 variable-lag pairs per grid point across all lags, of which 1.44 are contemporaneous ($\tau = 0$) and 1.25 are lagged ($\tau > 0$). Notably, more than half of the significant links are contemporaneous ($\tau = 0$) except for the PF and BKS-SIC; links at $\tau = 2, 3$ are essentially absent for all drivers, and the few that remain show no regional organisation.

The four drivers shown in Figure 2 each produce spatially coherent, dipole-like patterns of significant causal influence. SPV shows a dipole between Scandinavia (positive, rainier) and the Iberian Peninsula (negative, drier) under a strenghtened vortex; UB shows an opposing dipole with its negative center stretching from the Atlantic Ocean east of Newfoundland to Scandinavia and the Barents Sea, while the positive region comprises the Mediterranean and south-east Europe; BKS-SIC shows a strong negative link over Svalbard and its surroundings, together with a positive link over the Iberian Peninsula and to its east; finally, PF closely mirrors the classic NAO dipole.

The grid-point maps above provide a detailed, local picture. Since the algorithm is run independently at each grid point, with no information about neighbouring locations, the spatial organisation that is nevertheless visible suggests that precipitation in certain regions may be governed by the same set of large-scale drivers. To explore this more systematically, we apply k-means clustering to the grid points based on their causal driver vectors. This analysis is carried out using the causal link strengths shown in Figure A1. The spatial distribution of clusters is shown in Figure 3a, and the dominant large-scale drivers of each cluster are summarised in Figure 3c.

Cluster 1 is essentially driven by (PF, $\tau = 0$), with the centroid anticorrelated with precipitation. Weaker causal links also exist for (PF, $\tau = 1$), (SPV, $\tau = 0$), (UB, $\tau = 0$), and (BKS-SIC, $\tau = 0$). All of these show significant local correlations over the Mediterranean, Iberian Peninsula, and the Azores — precisely where the grid points of this cluster are located. Similar drivers but with opposite correlation sign govern cluster 2, located around the Icelandic Low and Scandinavia, again with (PF, $\tau = 0$) as the main driver. Cluster 3, by contrast, is mainly driven by (UB, $\tau = 0$), with its centre of action split between two subregions: east of Scandinavia and east of Newfoundland, consistent with the local analysis. Cluster 4, with grid points concentrated around Svalbard, is mainly driven by (BKS-SIC, $\tau = 0, 1$), again consistent with the local results. Finally, the centroid of cluster 0 shows no significant drivers, as this cluster groups together most of the Euro-Atlantic grid points for which no common set of large-scale drivers was identified. This does not mean that these grid points have no causal drivers: many of them have significant local drivers (Figure 2), but these differ from one grid point to another, so they do not group into a coherent structure in the space of causal driver vectors.

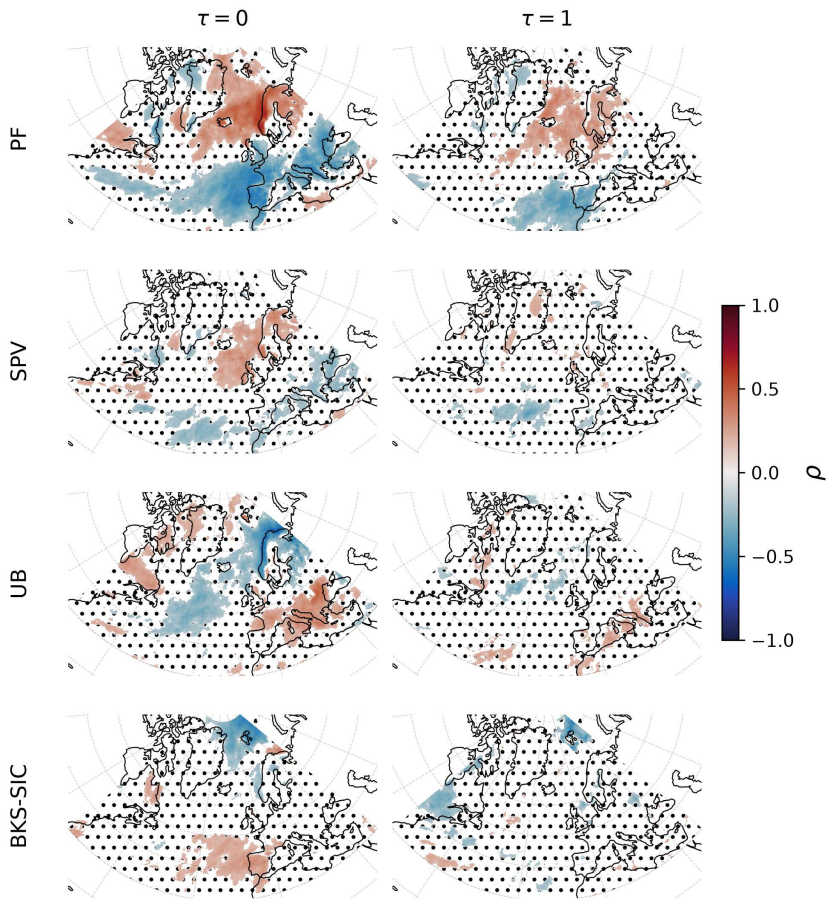


Fig. 2 Partial correlation between each variable-lag pair and DJF TP anomalies at every grid point, after conditioning on all other potential drivers retained by the PC stage of the algorithm. Rows correspond to large-scale drivers and columns to monthly lags. For example, the panel (UB, $\tau = 1$) shows the partial correlation between the Ural Blocking index during NDJ and precipitation during DJF, conditioned on all other potential confounders. Stippling indicates grid points where the result is not statistically significant ($p > 0.05$).

4.2 Regressing precipitation

Figure 4 compares the local regression skill (R^2) of the NAO against that of the causal drivers identified at each grid point. The causal drivers achieve significant regression skill over a substantially larger fraction of the domain than the NAO (88.81% vs. 67.92% of grid points). Local R^2 values remain modest in both cases, but the causal drivers extend meaningful regression skill into regions – away from the NAO's own centres of action – where the NAO itself carries little information. On average, the NAO achieves an R^2 of 0.16 over its significant grid points, compared with 0.15 for the causal drivers, so that the causal discovery algorithm improves the significant

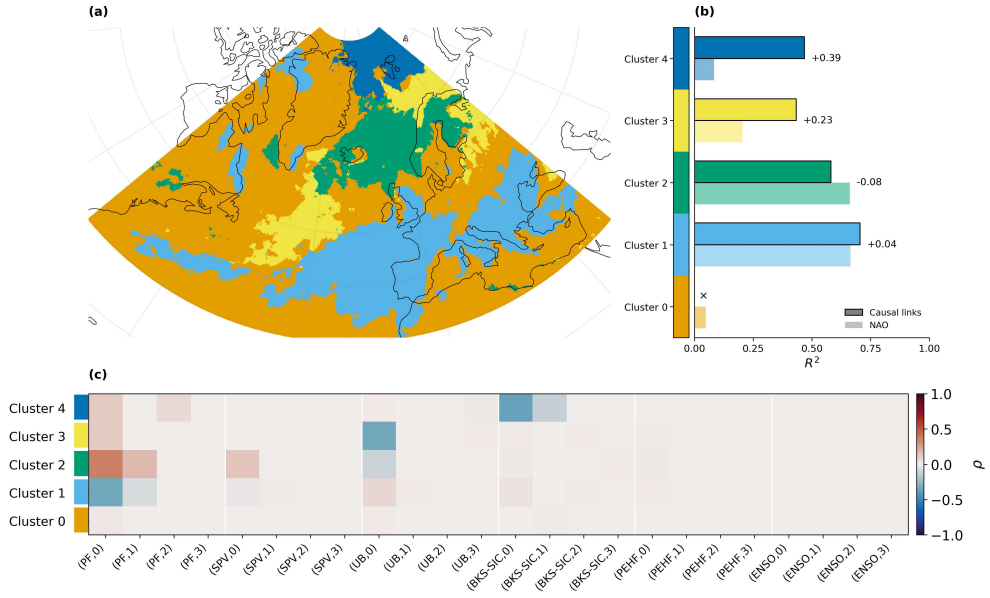


Fig. 3 a) Spatial distribution of the k-means clusters over the Euro-Atlantic domain. Each colour corresponds to a distinct cluster, grouping grid points with a similar causal structure. b) Regression skill (R^2) of the multiple linear regression of cluster-averaged DJF precipitation anomalies, using the cluster centroid drivers as regressors (darker bars), compared against the regression skill of the NAO index (lighter bars). Each colour corresponds to one cluster, as defined in a). The cross for cluster 0 indicates that no causal drivers were identified for that cluster, so no regression could be fitted. The numbers above each pair of bars show the difference $R^2_{\text{causal}} - R^2_{\text{NAO}}$, indicating how much the causal drivers improve upon the NAO benchmark. All regressions are significant ($p \leq 0.05$). c) Causal drivers of each cluster, represented by its centroid. Each row corresponds to a cluster (colour-coded on the left margin to match a)), and each column to a variable-lag pair. The colour of each cell encodes the partial correlation strength of each driver at the centroid of the cluster.

area by 20.89 percentage points over the NAO while leaving the average R^2 virtually unchanged.

To complement the local analysis, we regress the cluster-averaged precipitation, for the clusters identified in Figure 3a, onto the variable-lag pairs from the cluster centroids (Figure 3c). The cluster-averaged precipitation is also regressed onto the NAO index, providing a direct benchmark of the NAO's explanatory power at the regional scale. The resulting R^2 values are shown in Figure 3b, together with the difference $R^2_{\text{causal}} - R^2_{\text{NAO}}$, which indicates how much more variance the regression on the cluster-centroid drivers explains than the NAO benchmark. Regional averaging yields substantially higher R^2_{causal} values than the local analysis, of about 0.7 for clusters 1 and 2 and about 0.5 for clusters 3 and 4. For clusters 1 and 2, the explanatory power of the causal drivers is comparable to that of the NAO, whereas for clusters 3 and 4 the causal drivers outperform the NAO by 0.23 and 0.39 in R^2, respectively.

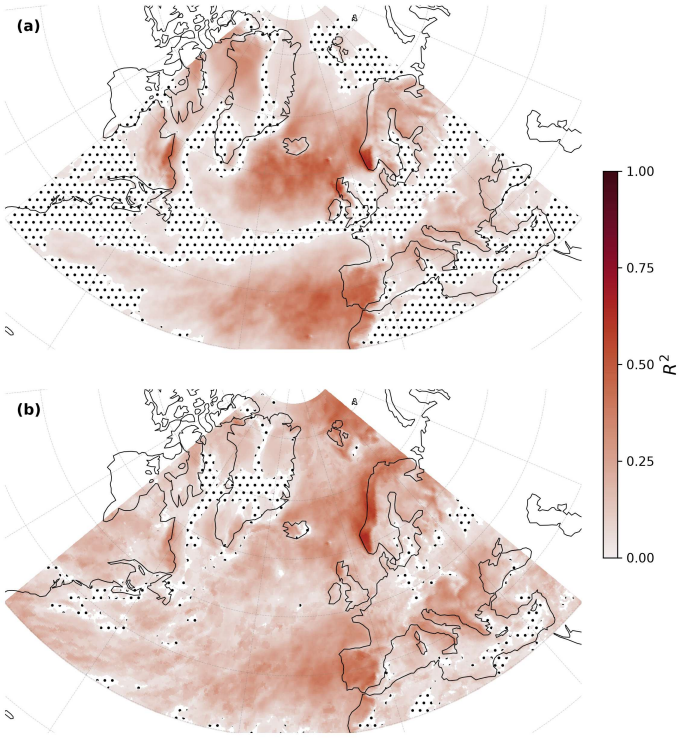


Fig. 4 Local regression skill (R^2) of the multiple linear regression of DJF precipitation anomalies at every grid point using as regression; the NAO index a) versus the drivers identified by the causality algorithm b). Stippling indicates grid points where the regression is not statistically significant ($p > 0.05$).

5 Discussion

A first observation regards the temporal structure of the links: at $\tau = 2, 3$ months few causal links exist, and those that do show no regional organisation, suggesting that the timescales of direct teleconnections are shorter than two months. We emphasise the word *direct*: as discussed in Section 2, signals from large-scale drivers reach precipitation primarily through messengers such as Rossby waves, which are short-lived and dynamic compared to our slowly varying drivers. In this framework, a direct link means that no other candidate driver lies on the pathway between the large-scale source and precipitation. This is consistent with significant direct links being confined to submonthly to monthly timescales, since longer-range influences, expected to be transmitted by intermediate drivers, would be attributed to those drivers rather than appear as direct links at long lags.

Turning to the individual drivers (Figure 2, and Figure A1), the results for ENSO and PEHF confirm their expected roles as an indirect driver and a mediator, respectively: ENSO may influence Euro-Atlantic precipitation, but indirectly, modulated by other large-scale drivers such as the PF (Brönnimann 2007; Hall et al. 2015), while the PEHF, a coupling mechanism between the troposphere and stratosphere, acts as

a messenger rather than as a driver in its own right, mediating the influence of the SPV on the troposphere (Hannachi et al. 2025). The dipole shown by the SPV at $\tau = 0$ is similar to that of the NAO, with positive and negative precipitation anomalies over Scandinavia and the Iberian Peninsula, respectively, under a strengthened vortex (Monnin et al. 2022; Zhang et al. 2022), but it is obtained here as a direct link, without invoking the NAO or the PF as intermediary mechanisms. The UB shows a dipole opposing the NAO at $\tau = 0$, consistent with a blocking-induced deflection of the storm track (Kautz et al. 2022; Murto et al. 2022). The BKS-SIC shows yet another dipole; the negative causal link with precipitation over Svalbard and its surroundings is consistent with the cold, high-pressure conditions associated with extensive sea ice (McKenna et al. 2018), and the positive influence on the Iberian Peninsula is in line with the previously reported connection between Arctic sea-ice loss and a reduced frequency of cut-off lows reaching southern Europe (Saurral et al. 2025). That the PF closely mirrors the NAO dipole at both $\tau = 0$ and $\tau = 1$ is unsurprising, given that the NAO pressure anomalies reflect the position of the PF; nonetheless, the PF adds a layer of physical interpretability, since its position tracks the storm track and hence the distribution of precipitation (Hall et al. 2015).

Most of the links discussed above are contemporaneous ($\tau = 0$), so their direction cannot be inferred from time ordering, in contrast with lagged links, for which a driver precedes precipitation by design. As explained in Section 3, contemporaneous links arise because our monthly resolution is coarser than some teleconnection timescales. We assign their direction on the physical argument that precipitation at a single grid point cannot drive a large-scale feature, so any such link is interpreted as an influence of the driver on precipitation operating on submonthly timescales. For instance, the BKS-SIC, an average over the entire Barents–Kara region, is not driven by precipitation at a single grid point, so the direction of the link follows from this reasoning.

Our causal discovery relies on linear conditional independence tests, which assume that the dependencies between the large-scale drivers and precipitation are well described by partial correlation. We also tested a non-linear variant based on distance correlation (Székely et al. 2007) with second-order polynomial residualisation, which is not a fully non-linear test. This variant did not provide additional or more robust links than the linear one, which we attribute to the short sample ($n \approx 130$) being insufficient to constrain more flexible dependencies. This does not rule out non-linear dependencies, and a fully non-linear framework, possibly supported by a larger sample such as that of climate model large ensembles, remains a possible direction for further analysis.

The five-cluster solution recovers spatially coherent regions even though the k-means algorithm has no access to geographical information, closely resembling the regions already suggested by the local analysis but condensing them into a single map. This provides a bridge between the local causal analysis, which identifies drivers at every grid point, and the identification of naturally emerging regions in which precipitation is governed by the same set of large-scale drivers, without the need to define ad hoc regions beforehand. Overall, the local and regional analyses show that a set of large-scale drivers identified through causal discovery can be found at most grid points across the Euro-Atlantic domain, and that these drivers do not appear

at random but organise into spatially coherent regions. This suggests that the NAO pattern can also be viewed as the summarised effect of several mechanisms acting contemporaneously on the North Atlantic and Europe, a view that offers added insight into the individual drivers and their possible future changes.

The above view is further supported by Figure 4, which constitutes the central result of this study, contrasting the local explanatory power of the NAO against that of the causally identified drivers. Figure 4a shows how limited the area explained by the NAO is, potentially a consequence of its constrained physical interpretability discussed above. When the causal drivers are instead used as regressors (Figure 4b), an immediate improvement is observed: while the local R^2 values remain modest in both cases, the significant area covers most of the Euro-Atlantic domain rather than being confined to the NAO's own centres of action. This improvement is partly explained by the fact that several causal drivers may be identified at each grid point, providing multiple regressors; we consider this an appropriate comparison rather than an unfair advantage, since these are not arbitrarily added regressors but the direct causal parents of precipitation identified by the algorithm. The generally low R^2 values across the domain reflect two factors: first, we rely exclusively on remote large-scale drivers rather than local thermodynamic variables, which also influence precipitation; second, local precipitation is inherently noisy and highly variable, which motivates the regional averaging that follows.

The higher R^2 values obtained after regional averaging in Figure 3b suggest that much of the variance left unexplained at individual grid points reflects the high local variability of precipitation, which averaging over a cluster smooths out. That the causal drivers of clusters 1 and 2 explain about the same variance as the NAO is expected, since each of these clusters approximately corresponds to one of the NAO's centres of action, and in both cases the regression is dominated by the same underlying driver, the PF. The added value of the causal framework is most apparent for clusters 3 and 4, where the larger explained variance reflects the ability of the UB and the BKS-SIC to account for precipitation in regions where the NAO carries little explanatory power. Together with Figure 4, these results support the central argument of this study: a physically grounded understanding of the mechanisms driving precipitation not only improves its interpretation, but also provides a better statistical description of it.

6 Conclusions

This study aimed to obtain a physically grounded characterisation of Euro-Atlantic winter precipitation by identifying, with a causal discovery framework, the direct large-scale drivers of its variability at every grid point. We identified such drivers across most of the domain, adding a layer of interpretability that the NAO, as an aggregate statistical descriptor, cannot provide by itself. As expected, no direct link was found for ENSO, while the SPV was found to be directly linked to precipitation across northern and southern Europe, and the BKS-SIC to Iberian precipitation. When clustering is performed on these drivers, they organise into spatially coherent regions that largely match the NAO's centres of action, supporting the view that the NAO summarises the

effect of several mechanisms. Some areas remain unresolved, most notably the transition zone between the two NAO poles, which future work could address by exploring additional candidate drivers such as sea surface temperatures, blocking patterns, or Eurasian snow cover.

This study also aimed to quantify the explanatory power of the identified causal drivers relative to the NAO. The central result of this paper, a causally-informed regression achieves significant regression skill over most of the Euro-Atlantic domain, extending well beyond the NAO's own centres of action, where the NAO itself carries little information. Local R^2 values remain modest, as expected given that we deliberately rely only on remote dynamical drivers rather than local thermodynamic variables; regionally, however, these same drivers demonstrate strong explanatory power even away from the NAO's poles, confirming that a physically grounded characterisation of precipitation translates into meaningfully improved statistical descriptions.

A promising next step would be to forecast the large-scale drivers themselves, whose smoother and more predictable evolution than local precipitation could provide an indirect route to improved precipitation forecasts, while also clarifying the full chain of indirect pathways that build up to drive precipitation beyond the direct links retained here. Expanding the training sample, for instance through climate model large ensembles or by moving to sub-monthly timescales, would further support more flexible regression frameworks than the linear approach adopted in this study.

Acknowledgements. RIS acknowledges his AI4S fellowship within the "Generación D" initiative by Res.es, Ministerio para la Transformación Digital y de la Función Pública of Spain, for talent attraction (C005/24-ED CV1), funded by NextGenerationEU through PRTR.

Declarations

Funding. JHC, JOM, and MPN acknowledge funding from Independent Research Fund Denmark through the project "Arctic push on Europe's extremes: The anatomy of Arctic impact on European climate change - ArcticPush" (Grant ID 10.46540/4258-00025B). MPN's research is funded by the Spanish grant PID2024-159949NA-I00 (MICIU/AEI/10.13039/501100011033).

Competing Interests. The authors have no relevant financial or non-financial interests to disclose.

Data availability. The data used in this analysis are all freely available. The ERA5 reanalysis (Hersbach et al. 2020) can be downloaded from the Copernicus Climate Data Store (https://cds.climate.copernicus.eu). The NAO index is available from the Climate Data Guide (Hurrel et al. 2026), and the Niño 3.4 index used as the ENSO index is available from the Climate Data Guide (Trenberth and Staff 2026).

Code availability. The code generated and used during the development of this study is available from the corresponding author on reasonable request.

Author Contributions. All authors contributed to the conceptualisation and design of the study. MPN performed data processing, statistical analysis, and methodological implementation. The initial draft of the manuscript was prepared by MPN and all authors contributed to subsequent revisions through critical review and intellectual input. All authors approved the final version of the manuscript.

References

Baba K, Shibata R, Sibuya M. Partial correlation and conditional correlation as measures of conditional independence. Australian & New Zealand Journal of Statistics. 2004;46(4):657–664. https://doi.org/10.1111/j.1467-842X.2004.00360.x.

Barnston AG, Chelliah M, Goldenberg SB. Documentation of a highly ENSO-related sst region in the equatorial pacific: Research note. Atmosphere-Ocean. 1997;35(3):367–383. https://doi.org/10.1080/07055900.1997.9649597.

Barnston AG, Livezey RE. Classification, Seasonality and Persistence of Low-Frequency Atmospheric Circulation Patterns. Monthly Weather Review. 1987;115(6):1083 – 1126. https://doi.org/10.1175/1520-493(1987)115⟨1083:CSAPOL⟩2.0.CO;2.

Berg P, Haerter JO, Thejll P, Piani C, Hagemann S, Christensen JH. Seasonal characteristics of the relationship between daily precipitation intensity and surface temperature. Journal of Geophysical Research: Atmospheres. 2009;114(D18). https://doi.org/https://doi.org/10.1029/2009JD012008.

Branstator G, Teng H. Tropospheric Waveguide Teleconnections and Their Seasonality. Journal of the Atmospheric Sciences. 2017;74(5):1513 – 1532. https://doi.org/10.1175/JAS-D-16-0305.1.

Brönnimann S. Impact of El Niño–Southern Oscillation on European climate. Reviews of Geophysics. 2007;45(3). https://doi.org/https://doi.org/10.1029/2006RG000199.

Caretta MA, Mukherji A, Arfanuzzaman M, Betts RA, Gelfan A, Hirabayashi Y, et al. Water. In: Pörtner HO, Roberts DC, Tignor M, Poloczanska ES, Mintenbeck K, Alegría A, et al., editors. Climate Change 2022: Impacts, Adaptation and Vulnerability. Contribution of Working Group II to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change. Cambridge, UK and New York, NY, USA: Cambridge University Press; 2022. p. 551–712.

Cohen J, Furtado JC, Jones J, Barlow M, Whittleston D, Entekhabi D. Linking Siberian Snow Cover to Precursors of Stratospheric Variability. Journal of Climate. 2014;27(14):5422 – 5432. https://doi.org/10.1175/JCLI-D-13-00779.1.

Cohen J, Screen JA, Furtado JC, Barlow M, Whittleston D, Coumou D, et al. Recent Arctic amplification and extreme mid-latitude weather. Nature geoscience. 2014;7(9):627–637. https://doi.org/10.1038/ngeo2234.

Ebert-Uphoff I, Deng Y. Causal Discovery for Climate Research Using Graphical Models. Journal of Climate. 2012;25(17):5648 – 5665. https://doi.org/10.1175/JCLI-D-11-00387.1.

Fisher RA. The distribution of the partial correlation coefficient. Metron. 1924;3:329–332.

Hall R, Erdélyi R, Hanna E, Jones JM, Scaife AA. Drivers of North Atlantic Polar Front jet stream variability. International Journal of Climatology. 2015;35(8). https://doi.org/10.1002/joc.4121.

Hanna E, Francis J, Wang M, Overland JE, Cohen J, Luo D, et al. Influence of high-latitude blocking and the northern stratospheric polar vortex on cold-air outbreaks under Arctic amplification of global warming. Environmental Research: Climate. 2024;3(4):042004. https://doi.org/10.1088/2752-5295/ad93f3.

Hannachi A, Straus DM, Franzke CLE, Corti S, Woollings T. Low-frequency nonlinearity and regime behavior in the Northern Hemisphere extratropical atmosphere. Reviews of Geophysics. 2017;55(1):199–234. https://doi.org/https://doi.org/10.1002/2015RG000509.

Hannachi A, Lechner M, Finke K, Mukhin D. Stratospheric polar vortex, wave absorption/reflection and effect on surface climate. Climate Dynamics. 2025;63(2):126. https://doi.org/10.1007/s00382-025-07610-1.

Hersbach H, Bell B, Berrisford P, Hirahara S, Horányi A, Muñoz-Sabater J, et al. The ERA5 global reanalysis. Quarterly Journal of the Royal Meteorological Society. 2020;146(730):1999–2049. https://doi.org/https://doi.org/10.1002/qj.3803.

Hoskins BJ, Karoly DJ. The Steady Linear Response of a Spherical Atmosphere to Thermal and Orographic Forcing. Journal of Atmospheric Sciences. 1981;38(6):1179 – 1196. https://doi.org/10.1175/1520-0469(1981)038⟨1179:TSLROA⟩2.0.CO;2.

Hurrel JW, Phillps A, Staff N.: The Climate Data Guide: Hurrell North Atlantic Oscillation (NAO) Index (PC-based); 2026. https://climatedataguide.ucar.edu/climate-data/hurrell-north-atlantic-oscillation-nao-index-pc-based.

Hurrell JW. Decadal Trends in the North Atlantic Oscillation: Regional Temperatures and Precipitation. Science. 1995;269(5224):676–679. https://doi.org/10.1126/science.269.5224.676.

Kautz LA, Martius O, Pfahl S, Pinto JG, Ramos AM, Sousa PM, et al. Atmospheric blocking and weather extremes over the Euro-Atlantic sector – a review. Weather and Climate Dynamics. 2022;3(1):305–336. https://wcd.copernicus.org/articles/3/305/2022/, https://doi.org/10.5194/wcd-3-305-2022.

Kim BM, Son SW, Min SK, Jeong JH, Kim SJ, Zhang X, et al. Weakening of the stratospheric polar vortex by Arctic sea-ice loss. Nature communications. 2014;5(1):4646. https://doi.org/10.1038/ncomms5646.

Kretschmer M, Adams SV, Arribas A, Prudden R, Robinson N, Saggioro E, et al. Quantifying Causal Pathways of Teleconnections. Bulletin of the American Meteorological Society. 2021;102(12):E2247 – E2263. https://doi.org/10.1175/BAMS-D-20-0117.1.

Kretschmer M, Coumou D, Donges JF, Runge J. Using Causal Effect Networks to Analyze Different Arctic Drivers of Midlatitude Winter Circulation. Journal of Climate. 2016;29(11):4069 – 4081. https://doi.org/10.1175/JCLI-D-15-0654.1.

McKenna CM, Bracegirdle TJ, Shuckburgh EF, Haynes PH, Joshi MM. Arctic Sea Ice Loss in Different Regions Leads to Contrasting Northern Hemisphere Impacts. Geophysical Research Letters. 2018;45(2):945–954. https://doi.org/https://doi.org/10.1002/2017GL076433.

Monnin E, Kretschmer M, Polichtchouk I. The role of the timing of sudden stratospheric warmings for precipitation and temperature anomalies in Europe. International Journal of Climatology. 2022;42(6):3448–3462. https://doi.org/https://doi.org/10.1002/joc.7426.

Murto S, Caballero R, Svensson G, Papritz L. Interaction between Atlantic cyclones and Eurasian atmospheric blocking drives wintertime warm extremes in the high Arctic. Weather and Climate Dynamics. 2022;3(1):21–44. https://doi.org/10.5194/wcd-3-21-2022.

Palmer TN. A Nonlinear Dynamical Perspective on Climate Prediction. Journal of Climate. 1999;12(2):575 – 591. https://doi.org/10.1175/1520-0442(1999)012⟨0575:ANDPOC⟩2.0.CO;2.

Pearl J. Causal inference in statistics: An overview. Statistics Surveys. 2009;3:96–146. https://doi.org/10.1214/09-SS057.

Peters J, Janzing D, Schölkopf B. Elements of causal inference: foundations and learning algorithms. Cambridge, MA: MIT Press; 2017.

Priestley MDK, Catto JL. Future changes in the extratropical storm tracks and cyclone intensity, wind speed, and structure. Weather and Climate Dynamics. 2022;3(1):337–360. https://doi.org/10.5194/wcd-3-337-2022.

Rogers JC. The Association between the North Atlantic Oscillation and the Southern Oscillation in the Northern Hemisphere. Monthly Weather Review. 1984;112(10):1999 – 2015. https://doi.org/10.1175/1520-0493(1984)112⟨1999:TABTNA⟩2.0.CO;2.

Runge J. Causal network reconstruction from time series: From theoretical assumptions to practical estimation. Chaos: An Interdisciplinary Journal of Nonlinear Science. 2018;28(7):075310. https://doi.org/10.1063/1.5025050.

Runge J, Bathiany S, Bollt E, Camps-Valls G, Coumou D, Deyle E, et al. Inferring causation from time series in Earth system sciences. Nature communications. 2019;10(1):2553. https://doi.org/10.1038/s41467-019-10105-3.

Runge J, Gerhardus A, Varando G, Eyring V, Camps-Valls G. Causal inference for time series. Nature Reviews Earth & Environment. 2023;4(7):487–505. https://doi.org/10.1038/s43017-023-00431-y.

Runge J, Nowack P, Kretschmer M, Flaxman S, Sejdinovic D. Detecting and quantifying causal associations in large nonlinear time series datasets. Science Advances. 2019;5(11):eaau4996. https://doi.org/10.1126/sciadv.aau4996.

Sanders F. Life History of Mobile Troughs in the Upper Westerlies. Monthly Weather Review. 1988;116(12):2629 – 2648. https://doi.org/10.1175/1520-0493(1988)116⟨2629:LHOMTI⟩2.0.CO;2.

Saurral RI, Doblas-Reyes FJ, Screen JA, Catto JL, Hay S, Yu H. Western Mediterranean Droughts Fostered by Arctic Sea Ice Loss. Journal of Climate. 2025;38(13):3005 – 3014. https://doi.org/10.1175/JCLI-D-25-0066.1.

Saurral RI, Kucharski F. Arctic amplification fingerprints on the midlatitude circulation in a medium-complexity global atmospheric model. Climate Dynamics. 2026;64(10):415. https://doi.org/10.1007/s00382-026-08373-z.

Shaw T, Baldwin M, Barnes EA, Caballero R, Garfinkel C, Hwang YT, et al. Storm track processes and the opposing influences of climate change. Nature Geoscience. 2016;9(9):656–664. https://doi.org/10.1038/ngeo2783.

Shepherd TG. Atmospheric circulation as a source of uncertainty in climate change projections. Nature Geoscience. 2014;7(10):703–708. https://doi.org/10.1038/ngeo2253.

Smith DM, Screen JA, Deser C, Cohen J, Fyfe JC, García-Serrano J, et al. The Polar Amplification Model Intercomparison Project (PAMIP) contribution to CMIP6: investigating the causes and consequences of polar amplification. Geoscientific Model Development. 2019;12(3):1139–1164. https://doi.org/10.5194/gmd-12-1139-2019.

Sneddon AK, Blamey RC, Reason CJC. Atmospheric River Climatology for the South Atlantic Basin Using an Image Processing-Based Methodology. Journal of Geophysical Research: Atmospheres. 2026;131(15):e2025JD045478. https://doi.org/https://doi.org/10.1029/2025JD045478.

Székely GJ, Rizzo ML, Bakirov NK. Measuring and testing dependence by correlation of distances. The Annals of Statistics. 2007;35(6):2769–2794. https://doi.org/10.1214/009053607000000505.

Tamarin-Brodsky T, Harnik N, Falkena SKJ. On Storm Tracks, Weather Regimes, and a Wave Breaking Recipe. AGU Advances. 2026;7(1):e2025AV002049. https://doi.org/https://doi.org/10.1029/2025AV002049.

Trenberth K, Staff N.: The Climate Data Guide: Nino SST Indices (Nino 1+2, 3, 3.4, 4; ONI and TNI); 2026. https://climatedataguide.ucar.edu/climate-data/nino-sst-indices-nino-12-3-34-4-oni-and-tni.

Vang A, Düsterhus A, Sørup HJD, Christensen JH. Decoding drivers of European precipitation: a multivariate mediation approach. Climate Dynamics. 2026;64(6):283. https://doi.org/10.1007/s00382-026-08232-x.

Wang L, Lu J, Breeden ML, Chen G, Henderson SA, Narinesingh V, et al. Gaps and ways forward in atmospheric blocking and extreme weather research. Nature Communications. 2026;17(1):2873. https://doi.org/s41467-026-70487-z.

Woollings T, Hannachi A, Hoskins B, Turner A. A Regime View of the North Atlantic Oscillation and Its Response to Anthropogenic Forcing. Journal of Climate. 2010;23(6):1291 – 1307. https://doi.org/10.1175/2009JCLI3087.1.

Ye K, Cohen J, Chen HW, Zhang S, Luo D, Hamouda ME. Attributing climate and weather extremes to Northern Hemisphere sea ice and terrestrial snow: progress, challenges and ways forward. npj Climate and Atmospheric Science. 2025;8(1):166. https://doi.org/10.1038/s41612-025-01012-0.

Zhang J, Zheng H, Xu M, Yin Q, Zhao S, Tian W, et al. Impacts of stratospheric polar vortex changes on wintertime precipitation over the northern hemisphere. Climate Dynamics. 2022;58(11):3155–3171. https://doi.org/10.1007/s00382-021-06088-x.

Appendix A Supplementary information

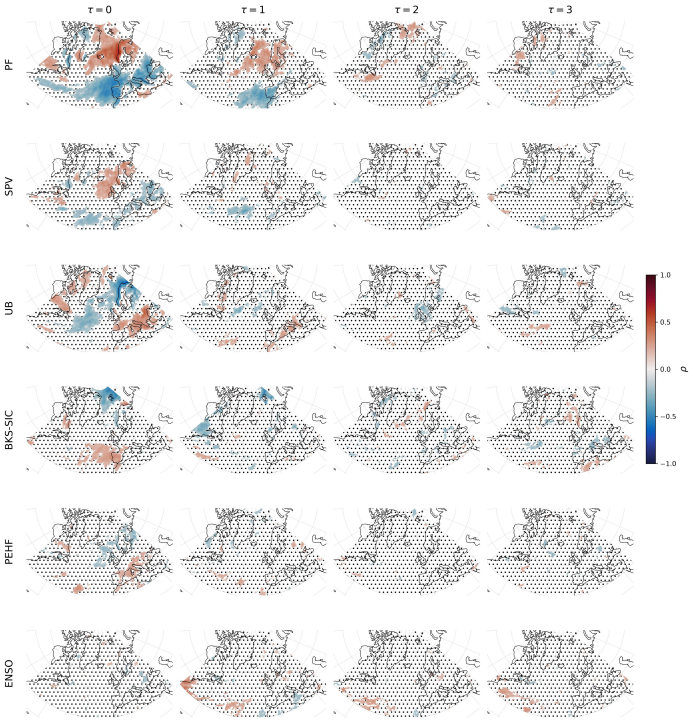


Fig. A1 Partial correlation between each variable-lag pair and DJF precipitation anomalies at every grid point, after conditioning on all other potential drivers retained by the PC stage of the algorithm. Rows correspond to large-scale drivers and columns to monthly lags. For example, the panel $(\mathrm{UB}, \tau = 2)$ shows the partial correlation between the Ural Blocking index during OND and precipitation during DJF, conditioned on all other potential confounders. Stippling indicates grid points where the result is not statistically significant ($p > 0.05$), assessed via a two-sided t-test with restricted degrees of freedom as explained in Section 3.