

ALMA MATER STUDIORUM · UNIVERSITY OF BOLOGNA

Department of Physics and Astronomy
Department of Biological, Geological and Environmental Sciences

Master Degree in Science of Climate

**A SOM-based analysis of the atmospheric
drivers of major European droughts: a
framework for extreme event attribution**

Supervisor:

Dr. Salvatore Pascale

Submitted by:

Amanda Montanaro

Co-supervisor:

Dr. Cristina Iacomino

Academic Year 2023/2024

Abstract

La siccità è uno dei fenomeni meteo-climatici più importanti. Negli ultimi decenni in Europa, siccità di estrema intensità e durata hanno avuto effetti deleteri per l'agricoltura, l'economia, gli ecosistemi e la società. Comprendere le cause di questi eventi è di fondamentale importanza, soprattutto alla luce del cambiamento climatico antropico, dal momento che l'intero continente, e in particolare la regione mediterranea, è riconosciuto come un'area critica del cambiamento climatico. In questo studio, abbiamo condotto un'analisi basata sulle Self-Organising Map (SOM) per classificare i dati mensili di altezza di geopotenziale a 500 hPa in ERA5 in prototipi di circolazione atmosferica dominanti sulla regione euro-atlantica. Queste configurazioni sono state poi analizzate in termini di variabilità atmosferica e di variabili legate alla siccità, valutando la loro relazione con le teleconnessioni euro-atlantiche ed il loro effetto sulle variabili idrologiche di superficie quali lo Standardized Precipitation Index (SPI) e lo Standardized Precipitation Evapotranspiration Index (SPEI). Utilizzando questo approccio, si sono analizzati tre eventi estremi di siccità europee di eccezionale intensità e durata: la siccità britannica del 1974-1975, la siccità iberica del 2004-2005 e la più recente siccità euromediterranea del 2022-2023. I nostri risultati mostrano che, durante questi eventi, la frequenza delle configurazioni di circolazione associate alle precipitazioni è stata di molto inferiore alla media, mentre quella delle configurazioni legate a condizioni secche molto superiore. Per esempio, nel caso della siccità del 2022-2023, la frequenza di configurazioni atmosferiche umide è stata del 3,7%, mentre quella di configurazioni secche ha raggiunto addirittura l'85,2%. Questa tecnica fornisce una metodologia che evidenzia i meccanismi atmosferici che governano le siccità ed il loro legame con la variabilità della circolazione su larga scala. Inoltre, consentendo una valutazione quantitativa della percentuale dei tipi di circolazione secca, umida e neutra, fornisce una metodologia per studi di attribuzione dell'effetto dei cambiamenti climatici antropogenici su eventi estremi di siccità.

Abstract

Drought is one of the most important meteo-climatic phenomena. In the last decades, European droughts of extreme intensity and duration have had severe impacts on agriculture, economy, ecosystems, and society. Understanding the drivers of these events has become increasingly crucial, especially in the context of anthropogenic climate change, as the whole continent, and, specifically, the Mediterranean region, is recognized as a climate change hotspot. In this study, we apply Self-Organizing Map (SOM) analysis to classify ERA5 500 hPa monthly geopotential height data into dominant circulation types over the Euro-Atlantic region. These patterns were then analyzed in terms of both atmospheric variability and drought-related variables, assessing their relation to Euro-Atlantic teleconnections and evaluating their effects on surface hydrological variables such as the Standardized Precipitation Index (SPI) and the Standardized Precipitation Evapotranspiration Index (SPEI). Using this framework, we examined three extreme European drought events characterized by extraordinary intensity and duration: the 1974-1975 UK drought, the 2004-2005 Iberian drought, and the most recent Euro-Mediterranean 2022-2023 drought. Our results show that, during these events, the frequency of precipitation-related circulation patterns significantly decreased, while that of dry patterns substantially increased. For instance, in the 2022-2023 case, the frequency of wet patterns dropped as low as 3.7%, whereas dry patterns reached up to 85.2%. This approach provides a methodology that highlights the atmospheric mechanisms driving droughts and their link to large-scale circulation variability. Most importantly, by allowing for a quantitative assessment of the percentage of the dry, wet, and neutral circulation types, it provides a framework for a future attribution studies on the effect of anthropogenic climate change on extreme droughts.

Contents

1	Introduction	5
1.1	Overview on droughts	5
1.2	Main metrics for drought assessment	7
1.3	Droughts in Europe and Italy	12
1.4	Relation between droughts in Europe and atmospheric variability	14
1.5	Self-Organising Maps (SOMs)	19
1.6	SOMs in drought studies	21
1.7	Objective of this thesis	23
2	Data and methods	24
2.1	Data	24
2.2	Technical details of the SOM algorithm	26
2.3	Methodology	30
2.3.1	Preprocessing of Z_{500}	31
2.3.2	Training of the 7×5 SOM	32
2.4	Composite analysis and statistical significance testing	34
2.5	Analysis of drought events	36
3	Results	39
3.1	SOM classification of atmospheric variability	41
3.1.1	Relationship of SOM nodes to main teleconnections	41
3.2	Relation between atmospheric patterns and drought-related quantities	48
3.3	SOM-based analysis of atmospheric drivers of drought events	52
3.3.1	The 1975-1976 drought	53

3.3.2 The 2004-2005 drought	60
3.3.3 The 2022-2023 drought	67
4 Discussion and Conclusions	74
4.1 Summary and discussion of the results	74
4.2 Conclusions	79
A Appendix	82

Chapter 1

Introduction

1.1 Overview on droughts

Drought, commonly characterized by below-normal precipitation over a period of months to years (Dai (2011b)), is one of the most impactful extreme meteorological events. Drought events, in fact, pose significant challenges due to their complex nature, making them difficult to monitor and predict. This complexity arises not only from the multiple environmental and climatic factors at play, such as rainfall patterns, soil evapotranspiration, river levels, and vegetation, but also from the wide range of temporal and spatial scales over which droughts may occur.

When characterized by long duration and large spatial extent, drought can result in considerable social, environmental and economic costs (Li et al. (2015), Spinoni et al. (2015a)) and can be seen as one of the most expensive and widespread natural disasters with negative impacts on agriculture, water resources, natural ecosystems, and society activities (Mishra and Singh (2010)). Droughts are among the most damaging natural disasters, causing tens of billions of dollars in damage and affecting millions of people all over the world each year (Wilhite (2000)). Furthermore, the exacerbation of drought events by global warming has increased the urgency to better un-

derstand and manage their impacts (Dai (2011a)).

The complexity of various influencing factors complicates the precise characterization of drought events, specifically their onset, duration, conclusion, extent, and, crucially, their underlying causes. Consequently, a universally accepted definition of drought is still missing. According to Wilhite and Glantz (1985), droughts are divided into 4 categories based on the type of water resource affected:

- **Meteorological drought:** it is defined as a lack of precipitation over a region for a period of time;
- **Agricultural drought:** it is associated with soil moisture deficit, which leads to reduced crop production and plant growth;
- **Hydrological drought:** it is defined as a deficit in surface and sub-surface water resources supply (e.g., rivers, lakes, etc);
- **Socioeconomic drought** occurs when the demand for an economic good exceeds supply as a result of a weather-related deficit in water supply.

Meteorological droughts refer to instances in which precipitation deficits last between 1 and 3 months. However, critical shortages in soil moisture and water reservoirs, common characteristics of the other drought categories, may require longer periods (from several months to years) of lack of precipitation. This suggests that the period considered is essential to differentiate between the various forms of drought (Pascale and Ragone (2025)).

Droughts lasting more than a year, also known as multi-year droughts, can severely compromise water security, leading to irreversible socioeconomic and ecological consequences. One notable example of a multi-year drought is the "Day Zero" drought in Cape Town, which lasted from 2015 to 2018. This exceptional multi-year event resulted in one of the worst water crises in the metropolitan area and has been widely studied in scientific literature (e.g. Odoulami et al. (2020, 2023), Pascale et al. (2020)).

1.2 Main metrics for drought assessment

The assessment of the sensitivity of a drought event and real-time monitoring requires the use of quantitative indices. The Standardized Precipitation Index, or SPI (McKee et al. (1993)), the Standardized Precipitation-Evapotranspiration Index, or SPEI (Vicente-Serrano et al. (2010)), and the Palmer Drought Severity Index, or PDSI (Palmer (1965)), are among the most widely used for this aim, due to the simplicity of their definition and to their flexibility to describe droughts at multiple time scales. The SPI, introduced to measure precipitation deficits, offers temporal flexibility by allowing analysis over various timescales; however, it focuses exclusively on precipitation, potentially limiting its ability to capture water balance dynamics. The definition of the SPI (Bordi and Sutera (2002)), is based on the assumption that the cumulative precipitation over a period (e.g. 3,12,24 months) is a random variable that is gamma-distributed. Thereby, if x is the random variable, the gamma distribution is defined by its probability density function as

$$g(x) = \frac{1}{\beta^\alpha \Gamma(\alpha)} x^{\alpha-1} e^{-\frac{x}{\beta}} \quad (1.1)$$

where $\alpha > 0$ is a shape parameter, $\beta > 0$ is a scale parameter and $\Gamma(\alpha)$ is the gamma function. These parameters, which can be estimated from the data sample by means of a maximum-likelihood method, are then used to find the cumulative probability of precipitation for a given month and time scale at the station considered. The cumulative probability, letting $t = x/\beta$, becomes the incomplete gamma-function:

$$G(x) = \int_0^x g(x) dx = \frac{1}{\Gamma(\alpha)} \int_0^x t^{\alpha-1} e^{-t} dt \quad (1.2)$$

Since the gamma function is undefined for $x = 0$ and the precipitation field may contain zeros, the cumulative probability becomes:

$$H(x) = q + (1 - q)G(x) \quad (1.3)$$

where q is the probability of zero precipitation. $H(x)$ is then transformed to a normal variable Z through the following approximation:

$$\begin{aligned} Z = SPI &= -(t - \frac{c_0 + c_1 t + c_2 t^2}{1 + d_1 t + d_2 t^2 + d_3 t^3}) \quad \text{for } 0 < H(x) \leq 0.5, \\ Z = SPI &= +(t - \frac{c_0 + c_1 t + c_2 t^2}{1 + d_1 t + d_2 t^2 + d_3 t^3}) \quad \text{for } 0.5 < H(x) < 1 \end{aligned} \quad (1.4)$$

where

$$\begin{aligned} t &= \sqrt{\ln(\frac{1}{H(x)^2})} \quad \text{for } 0 < H(x) \leq 0.5, \\ t &= \sqrt{\ln(\frac{1}{1 - H(x)^2})} \quad \text{for } 0.5 < H(x) < 1 \end{aligned} \quad (1.5)$$

and $c_0, c_1, c_2, d_1, d_2, d_3$ are the following constants:

$$\begin{aligned} c_0 &= 2.515517, & d_1 &= 1.432788, \\ c_1 &= 0.802853, & d_2 &= 0.189269, \\ c_2 &= 0.010328, & d_3 &= 0.001308. \end{aligned} \quad (1.6)$$

The definition of the SPEI (Vicente-Serrano et al. (2010)) is similar to that of the SPI, but it expands on this by incorporating potential evapotranspiration (PET). As such, it measures the accumulated deficit of precipitation - evapotranspiration, i.e. the climatic water balance at the surface. The SPEI is particularly useful in regions where temperature-driven evapotranspiration plays a significant role in water scarcity, and it allows for a clear differentiation between meteorological drivers of a drought, associated with precipitation, and thermodynamic effects linked to temperature. In fact, to a first approximation, PET depends on temperature alone. Thornthwaite equation (Thornthwaite (1951)) estimates PET based primarily on temperature and latitude, making it relatively simple but less accurate in arid re-

gions. The Penman-Monteith equation (Monteith (1965)) incorporates other variables such as humidity, wind speed, and solar radiation, which also determines PET, thus providing a more accurate estimate. Finally, the Hargreaves equation (Hargreaves and Samani (1985)) offers a simpler alternative to Penman-Monteith, relying mainly on temperature and extraterrestrial radiation, making it useful when limited data are available.

Finally, PDSI, as with SPEI, incorporates both PET and precipitation data to assess long-term drought severity, although its main limitation is its fixed temporal scale and reliance on simplified soil moisture models. Some of the PDSI problems were solved with the development of the self-calibrated PDSI (sc-PDSI) (Wells et al. (2004)), although the main shortcoming of the PDSI remains. The choice of the drought index often depends on the specific objectives of the analysis and the temporal or spatial scales of interest. For example, when evaluating meteorological droughts, SPI-3 (or SPEI-3) is commonly used. The delay in the transition from meteorological to hydrological drought suggests that SPI-12 (SPEI-12) may be appropriate for the types of events (Van der Wiel et al. (2022)) since they are less sensitive to seasonal variations and singular extreme events (McKee et al. (1993); Spinoni et al. (2015a)). Nevertheless, if the objective is to highlight seasonality or the interplay between droughts and particular seasonal atmospheric circulation patterns, as explored by Kingston et al. (2015), employing SPI-6 (SPEI-6) may bring significant insights.

Fig 1.1 shows that SPI and PDSI are almost equivalent over long periods (i.e., 12 months, 24 months). Furthermore, the comparison of SPI and SPEI allows to determine their reliability and to explore the diversity and complexities of the various factors that affect the phenomenon in question. For instance, different values for SPI and SPEI might underline the different roles played by soil moisture and precipitation. For example, Baronetti et al. (2020) focused on this difference to separate the driving factor for drought in the Po Plain.

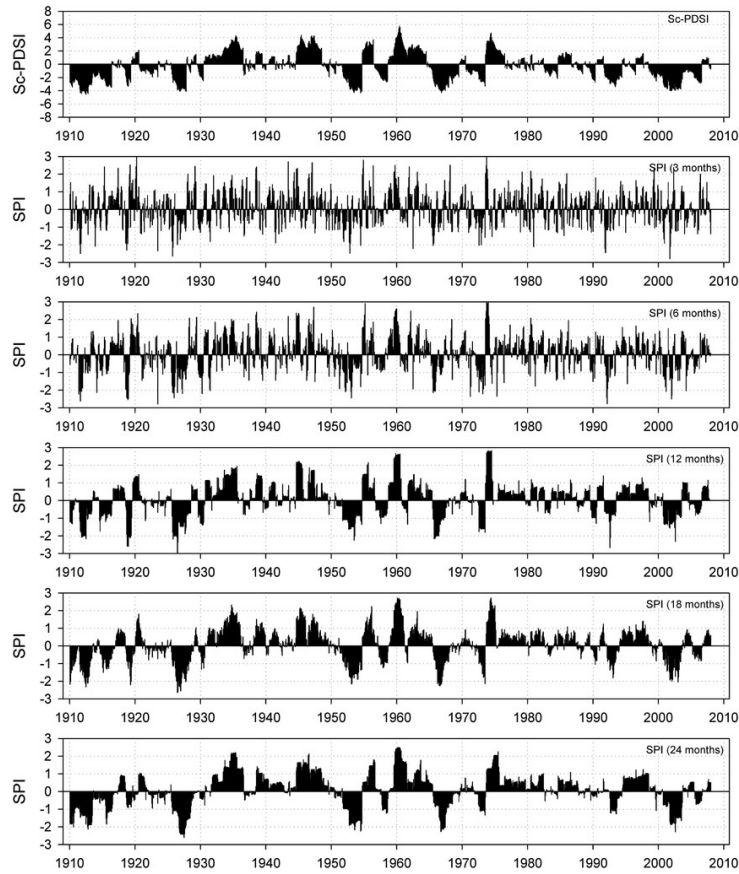


Figure 1.1: The sc-PDSI and 3-, 6-, 12-, 18-, and 24-month SPI at the Indore observatory (1910–2007) (Vicente-Serrano et al. (2010)).

For this study, SPI and SPEI were identified as the most suitable indices, as they can be easily worked out from observations and climate model outputs. Both SPI and SPEI are constructed as cumulative sums over a chosen number of months, allowing a temporal flexibility which is fundamental for drought analysis, given the wide range of scale of the phenomenon. The indexes most frequently employed vary according to the category of interest, with SPI-1, SPI-3, SPI-6, SPI-12, and SPI-24 (alongside SPEI-1, SPEI-3, SPEI-12, and SPEI-24) being predominant.

SPI and SPEI values typically range from -3 to 3. The severity of a drought

event can be characterized by the thresholds given in Table 1.1, with values less than -1 being commonly used to define drought conditions. The index

SPI/SPEI Value	Dry/wet Category
-2.0 or less	Extreme drought
-1.5 to -1.99	Severe drought
-1.0 to -1.49	Moderate drought
-0.99 to 0.99	Near normal
1.0 to 1.49	Moderate wet
1.5 to 1.99	Severe wet
2.0 or more	Extreme wet

Table 1.1: Table of SPI (SPEI) thresholds corresponding to different dry/wet categories (McKee et al. (1993)).

computation generally occurs at grid points or at a given meteorological station. Regions of the domain exhibiting a common SPI (SPEI) value below a specified threshold are identified as drought-affected areas. These kinds of indices are frequently employed to track the status of drought occurrences, but also to establish correlations with climate indices that identify both atmospheric and non-atmospheric modes. Fig. 1.2 shows drought conditions in Europe in terms of SPEI-12 in December 2024: Eastern Europe, the Balkan area, and the Mediterranean region seem to be affected by low values of SPEI-12, indicating low levels of precipitation at least in the last 12 months. On the other hand, the northwest region of Europe has positive values of SPEI, indicating the presence of weather regimes enriched in precipitation in the last 12 months, especially in the UK.

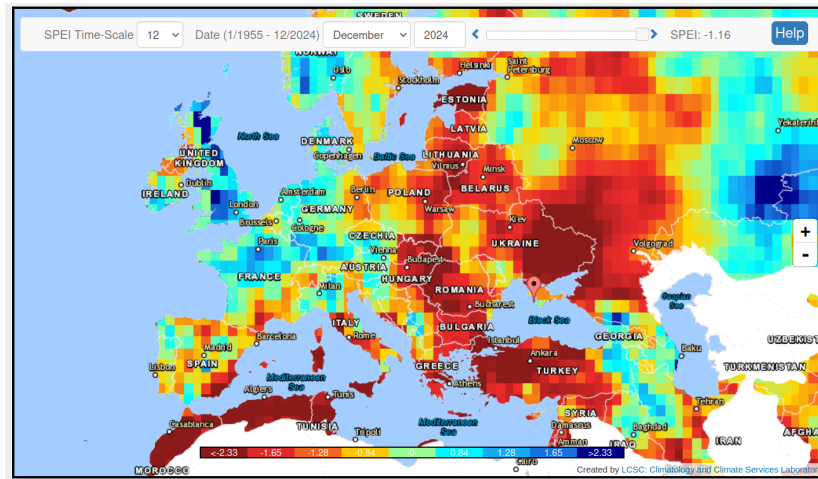


Figure 1.2: Standardized Precipitation Evapotranspiration Index (SPEI-12) for the 12-month accumulation period ending in December 2024, from the SPEI Global Drought Monitor (<https://spei.csic.es/map/maps.html#months=4#month=11#year=2024>).

1.3 Droughts in Europe and Italy

Droughts in Europe can be identified on different scales, both temporally and spatially, ranging from short-term, localized events to prolonged, widespread conditions that span multiple seasons or even years. Several studies show that, in the last 100 years, major droughts have affected not only areas such as southern Europe (Garrido-Perez et al. (2024)), but also western and central Europe, the British Islands (Parry et al. (2012)), Scandinavia, eastern Europe (Spinoni et al. (2015b)), and Russia, as discussed in Parry et al. (2012).

Past and current drought conditions are monitored and assessed by national and international agencies, such as the European Drought Observatory, or EDO, (<http://edo.jrc.ec.europa.eu/>) and the SPEI Global Drought Monitor (<http://sac.csic.es/spei/>), actively track and assess past and current drought conditions, providing critical data for mitigation and adap-

tation efforts. An example of drought assessment for Europe can be seen in Fig 1.3, showing, for July 2024, wetter-than-normal conditions from the Po Plain to the Nordic Sea, and drier-than-normal conditions in the region between the Black Sea and the Caspian Sea.

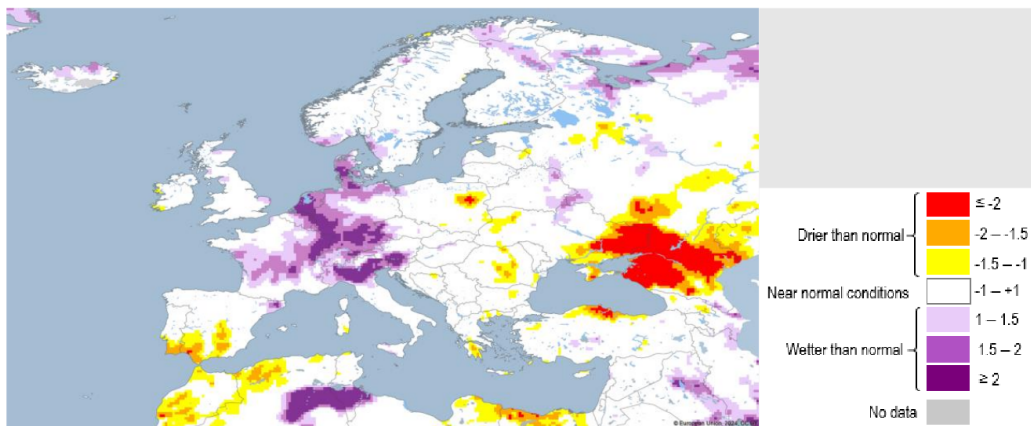


Figure 1.3: SPI-3 in July 2024, from the EGDO Analytical Report Global Drought Observatory: Drought in Europe (<https://drought.emergency.copernicus.eu/>).

Since the second half of the 20th century, Europe has experienced a slightly increasing drought pattern (Dai (2011a); Spinoni et al. (2015a); García-Herrera et al. (2007)) partly due to global warming. Rising temperatures caused by the effect of global warming are contributing to an increase in the atmosphere’s capacity to hold water vapor. This thermodynamic relation is well explained by the Claius-Clapeyron equation, with supporting evidence illustrated in the Assessment Report 6 (AR6), WG1 (IPCC (2023)). As the atmospheric demand for water vapor rises, there is a corresponding increase in potential evapotranspiration, which leads to increased drying of the soil. This effect is particularly evident in southern Europe (Spinoni et al. (2015a), Faranda et al. (2023)). Identified as a climate change hotspot (Giorgi (2006)),

the Mediterranean region is experiencing more frequent and intense droughts, often linked to rising temperatures and record-breaking heatwaves (Garrido-Perez et al. (2024)).

Regarding Italy, different approaches by different authors have been used to assess drought conditions in the peninsula. For instance, Baronetti et al. (2020) investigated drought occurrences in the Po Plain over the past fifty years, examining their underlying causes and their relationship with climate change (Baronetti et al. (2022)). In a broader context, Bordi and Sutera (2002) explored the influence of low-frequency atmospheric variability on drought through SPI. In particular cases, emphasis was placed on specific events: Faranda et al. (2023) examined the 2022 drought event, which had repercussions in northern Italy, southern France, and Spain. More recently, Pascale and Ragone (2025) compiled an extensive list of significant widespread multi-year drought events that affected the whole country from 1901 to the present day. Using SPI-12 and SPEI-12 they identified 9 major events, spanning from the drought of 1921-1922 and extending to the most recent event from 2021 to 2023. Other significant drought periods were recorded years 1942-1944, 1945-1946, 1979-1981, 1988-1989, 2006-2008, 2011-2013, and 2015-2019.

1.4 Relation between droughts in Europe and atmospheric variability

The main low-frequency variability modes, or simply, teleconnections, play an important role in driving drought events in Europe, as below-average precipitation is generally associated with persistent anticyclonic conditions related to specific large-scale atmospheric patterns. The main low-frequency teleconnections over the Euro-Atlantic sector were first identified in Barnston and Livezey (1987) through EOF analysis. As an example, Fig. 1.4 shows

the first 4 EOFs of Z_{500} anomalies (in order of explained variance), which physically correspond to: a) the North Atlantic Oscillation (NAO); b) the Scandinavian Pattern (SCAND); c) the East Atlantic (EA); d) the East Atlantic/Western Russian (EAWR). The prominence of the spatial pattern of these modes varies seasonally and it is particularly evident in winter. When differentiating between an extended winter (December through May) and an extended summer (June to November), the percentage of explained variance increases during the winter months and decreases during the summer months. This can be seen in Fig [A.1](#) and Fig [A.2](#) for winter and summer, respectively.

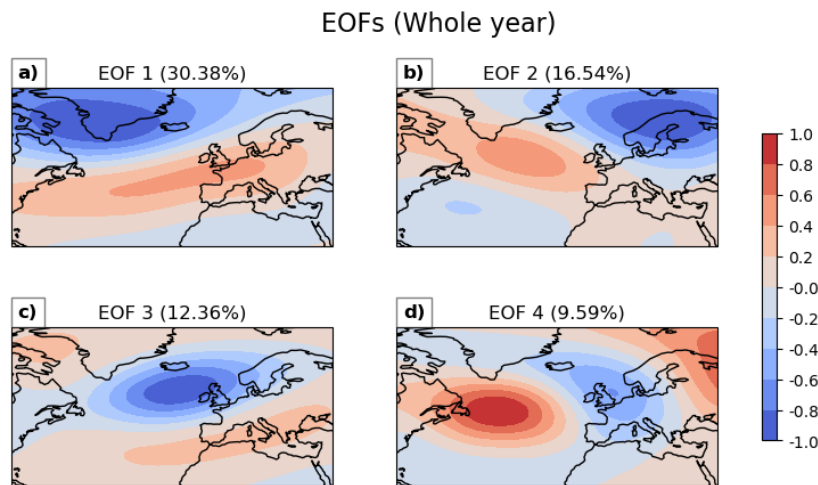


Figure 1.4: EOF loading pattern of monthly Z_{500} anomalies of: a) EOF1 (NAO), b) EOF2 (SCAND), c) EOF3 (EA), d) EOF4 (EAWR). Based on ERA5 monthly reanalysis ([Hersbach et al. \(2020\)](#)). Percentage over each panel denotes the fraction of explained total variance.

According to the NOAA Climate Prediction Center (<https://www.cpc.ncep.noaa.gov/data/teledoc/telecontents.shtml>), these four main teleconnection patterns are defined as follows:

- The **NAO** (Fig 1.4.a) consists of a north-south dipole of anomalies, with one center located over Greenland and the other center of opposite sign spanning the central latitudes of the North Atlantic between 35 °N and 40 °N. The positive phase of the NAO reflects below-normal heights and pressure across the high latitudes of the North Atlantic and above-normal heights and pressure over the central North Atlantic, the eastern United States, and western Europe. The negative phase reflects an opposite pattern of height and pressure anomalies over these regions.. The 3-month running mean (Fig. 1.5) shows a substantial variability at the seasonal and multi-annual timescale.
- The **SCAND** (Fig 1.4.b) consists of a primary circulation center over Scandinavia, with weaker centers of opposite sign over western Europe and eastern Russia / western Mongolia. The Scandinavian pattern has previously been referred to as the Eurasia-1 pattern by Barnston and Livezey (1987). The positive phase of this pattern is associated with positive height anomalies, sometimes reflecting major blocking anticyclones, over Scandinavia and western Russia, while the negative phase of the pattern is associated with negative height anomalies in these regions.
- The **EA** pattern (Fig 1.4.c) is the second prominent mode of low-frequency variability over the North Atlantic, and appears as a leading mode in all months. The EA pattern is structurally similar to the NAO, and consists of a north-south dipole of anomaly centers spanning the North Atlantic from east to west. The anomaly centers of the EA pattern are displaced southeastward to the approximate nodal lines of the NAO pattern. For this reason, the EA pattern is often interpreted as

a 'southward shifted' NAO pattern. However, the lower-latitude center contains a strong subtropical link in association with modulations in the subtropical ridge intensity and location.

- The **EAWR** (Fig 1.4.d) pattern is one of three prominent teleconnection patterns that affect Eurasia throughout the year. This pattern has been referred to as the Eurasia-2 pattern by Barnston and Livezey (1987). The East Atlantic / West Russia pattern consists of four main anomaly centers. The positive phase is associated with positive height anomalies located over Europe and northern China and negative height anomalies located over the central North Atlantic and north of the Caspian Sea.

The relationship between atmospheric teleconnections and European droughts has been extensively studied in scientific literature. For example, Vicente-Serrano and López-Moreno (2008) investigated the European response of the SPI-12 to the NAO and found opposing NAO–SPI correlations between northern and southern Europe (positive correlation in the north; negative correlation in the south). Furthermore, while Bordi and Sutera (2002) identified the NAO as a key driver for drought in Europe, other studies such as Sousa et al. (2011) and Pascale and Ragone (2025) support the hypothesis of the NAO and the SCAND being the primary factors influencing droughts in the western and central Mediterranean regions. In addition, Parry et al. (2012) and Kingston et al. (2015) explored the effects of other circulation patterns such as the EAWR patterns on famous drought events in the UK. However, there remains uncertainty over the connections in the process cascade that link such large-scale low-frequency variability to precipitation deficits. A given phase of the NAO, or any other atmospheric modes, does not always directly lead to the expected precipitation anomaly in a specific region, and drought indices are moderately correlated or not correlated at all (Kingston et al. (2015)). This was the case, for example, during the 2004-2005 Iberian drought (García-Herrera et al. (2007)), when atmospheric blocking condi-

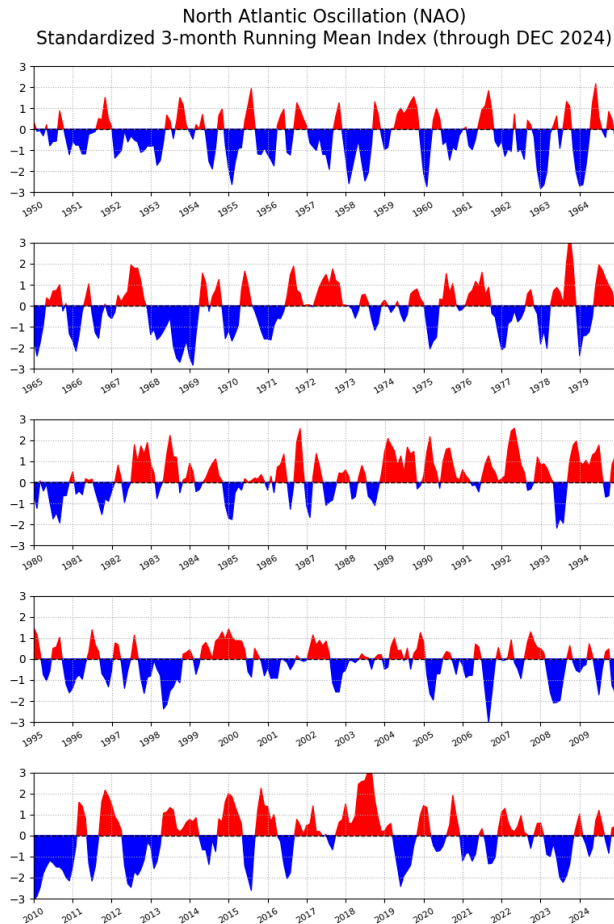


Figure 1.5: NAO timeseries from the National Oceanic and Atmospheric Administration (NOAA) (<https://www.cpc.ncep.noaa.gov/data/teledoc/telecontents.shtml>). The index is based on the sea-level pressure difference between the Subtropical (Azores) High and the Subpolar Low.

tions over the central Atlantic observed in certain months (e.g. March 2005) and associated with below-average precipitation, occurred with NAO and EA phases opposite of what typically expected with droughts over Iberia (specifically, the NAO was in its negative phase, while the EA in its positive phase). This further supports the argument that atmospheric modes alone

may not be sufficient to study the relationship between droughts and atmospheric variability, as their relationship with precipitation anomalies can be highly nonlinear and modulated by additional factors. Reflecting this complexity, connections between circulation indices, local climate, and hydrological variation (including drought) are often more complex than can be described adequately through index-based approaches—as demonstrated by analyses based on gridded climate fields (Parry et al. (2012); Kingston et al. (2015)). This suggests that simply performing a correlation analysis between atmospheric modes and drought may be insufficient. A more flexible approach to characterize variability of the atmospheric state is required, especially for intermediate configurations that do not distinctly match the positive or negative phases of known atmospheric modes. Self-organizing maps may provide such approach.

1.5 Self-Organising Maps (SOMs)

Self-Organizing Maps (SOMs), sometimes also referred to as Kohonen maps, were first introduced by Ritter and Kohonen (1989). A SOM is a type of artificial neural network (ANN) trained using unsupervised learning to generate a low-dimensional (typically two-dimensional) representation of multi-dimensional input data. This representation, referred to as a self-organized map, reduces the dimensionality of large datasets by clustering similar data points and organizing them into a structured two-dimensional grid (Skific and Francis (2012)). In many aspects, SOMs are analogous to other forms of cluster analysis; in fact, given an N-dimensional cloud of data points, the SOM seeks to place an arbitrary number of nodes within the data space such that their distribution represents the multi-dimensional density function. The nodes are spaced more closely in regions of high data density. However, SOMs differ from traditional cluster algorithms in one significant way, that is how groups are defined. While the end result of the SOM analy-

sis is some form of data clustering, SOMs do not directly partition data into groups. They create a map where each node represents a nearby set of data points. While clusters may form naturally, the primary goal of SOMs is to represent the overall data structure rather than explicitly defining clusters. SOMs attempt to find nodes or points in the measurement space that are representative of the nearby cloud of observations and, when taken together, describe the multi-dimensional distribution function of the data set (Hewitson and Crane (2002)).

SOMs in climate science and meteorology have been increasingly used in the last two decades. Some of the most notable applications are synoptic climatology (Hewitson and Crane (2002)) as well as North Atlantic variability (Reusch et al. (2005)), meteorological extremes (Cavazos (2000)), ENSO flavor distinction (Johnson (2013)) and NAO shift analysis, particularly regarding its secular eastward shift that began in the 1970s (Johnson et al. (2008)). SOMs provide a more flexible, robust, and interpretable framework for analyzing climate data than other more traditional techniques, such as EOFs. Unlike the latter, which by construction decompose data into orthogonal modes of variability based on linear assumptions, SOMs accommodate non-linear relationships in the data, making them better suited for capturing complex patterns and processes that are often present in atmospheric and oceanic systems.

One key advantage of SOMs over EOFs lies in their ability to represent data as a continuum of states rather than discrete modes. This characteristic is particularly valuable for atmospheric circulation data, where SOMs capture the full spectrum of synoptic conditions. In contrast, EOFs impose a rigid framework by representing variability as mutually orthogonal modes, potentially oversimplifying the underlying dynamics. This was proven to be extremely helpful in the discussion of the eastward shift to the NAO (Johnson et al. (2008)), and more in general, to the whole context of the North Atlantic large-scale circulation (Rousi et al. (2015)), where SOMs demon-

strated their ability to uncover subtle shifts and transitions that are difficult to capture with EOFs (Reusch et al. (2007)).

1.6 SOMs in drought studies

Recently, SOMs have proven to be useful in attribution studies for extreme drought events, i.e., in determining the effect of anthropogenic climate change (ACC) and the likelihood of circulation types driving droughts (Otto, 2017).




Figure 1.6: A 3×4 SOM for New Zealand 2013 event. Inset numbers refer to the number of days each node was observed over the duration of the drought from Harrington et al. (2016).

A notable application of SOMs for this purpose is found in the study by Harrington et al. (2016), which analyzed the 2013 New Zealand drought (Fig. 1.6). They used SOMs within a multi-member coupled climate model ensemble to examine changes in daily circulation patterns between two 41-year periods: a counterfactual climate with anthropogenic forcing as in 1861-1991

and a factual climate with anthropogenic forcing as in 1993–2023. Their results suggest a strong anthropogenic influence on the likelihood of exceptionally high seasonal mean sea level pressure patterns over New Zealand, with the most extreme manifestation occurring during the 2013 drought. In particular, they observed an increase in the frequency of nodes corresponding to blocking configurations.

Another significant example is the study by [Odoulami et al. \(2020\)](#) and later by [Odoulami et al. \(2023\)](#), who applied SOMs to analyze the 2015–2018 Western Cape drought in South Africa (Fig. [1.7](#)). Knowing the frequency of occurrence of each node, they were able to assess the persistence of specific configurations and quantify their associated rainfall. Using this quantity as reference, they classified "dry" and "wet" nodes and identified two key contributors to the event: a decrease in the frequency of rain-producing circulation patterns, accompanied by an increase in the occurrence of "dry" nodes, and a reduction in the rainfall volume associated with "wet" nodes.

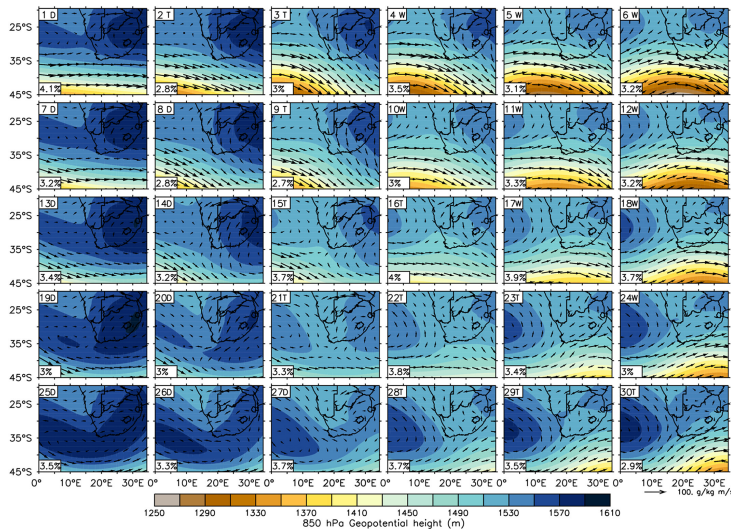


Figure 1.7: 6×5 SOM by [Odoulami et al. \(2020\)](#) for an attribution study in the "Day Zero" Cape Town drought of 2015-2018.

Expanding on this analysis, [Odoulami et al. \(2023\)](#) linked the event to human-induced climate change, using five large ensembles of atmospheric and coupled model simulations to compare climate scenarios with natural-only (i.e., the counterfactual climate) and anthropogenic (i.e., the factual climate) forcings.

1.7 Objective of this thesis

The objective of this thesis is to develop an attribution study framework for major European, multi-year droughts based on SOMs. The SOM algorithm will be applied to atmospheric reanalysis to identify distinct anomaly patterns in atmospheric circulation. These configurations will then be examined in terms of their relationship with low-frequency variability, assessing their connection to the main atmospheric teleconnections and their significance for variables associated with the surface climatic water balance such as precipitation, temperature, and potential evapotranspiration. This analysis will highlight how different low-frequency atmospheric configurations influence drought patterns across the study domain. This framework will be applied to three major prolonged drought events: the 1974–1975 UK drought, the 2004–2005 Iberian drought, and the 2022–2023 Euro-Mediterranean drought. This methodology will allow us to analyze the frequency and persistence of SOM nodes corresponding to dry, wet and neutral circulation patterns, and to provide a deeper understanding of how persistence in circulation patterns contributes to drought events. Finally, it will be explained how this framework can be used for extreme event attribution studies to assess the influence of anthropogenic climate change on the frequency of occurrence of these dry, wet and neutral circulation types during a drought event.

Chapter 2

Data and methods

In this chapter, we describe the dataset used to characterize low-frequency patterns through the SOMs and provide additional technical details about SOMs. In particular, we will provide additional technical details regarding SOMs. Furthermore, we outline the training and refinement processes involved in developing the model, highlighting the key steps taken to ensure optimal performance. In addition, we discuss essential preprocessing techniques and statistical methods that play a crucial role in analyzing the results. Finally, we discuss the major steps for the analysis of three relevant drought events. These events were selected following a comprehensive review of the literature on major droughts in Europe, identifying them as particularly impactful cases in terms of severity (more details in Section 3.3).

2.1 Data

To analyze the low-frequency atmospheric variability in Europe and the North Atlantic Ocean (spatial domain defined by [80°W, 40°E, 20°N, 80°N]) we use the geopotential height at 500 hPa (Z_{500}). Z_{500} data were obtained from the ERA5 reanalysis (Hersbach et al. (2020)) at monthly frequency. ERA5 is the fifth generation ECMWF atmospheric reanalysis of the global

climate covering the period from January 1940 to the present and it is produced by the Copernicus Climate Change Service (C3S) at ECMWF. Re-analyses combine past observations with models to generate consistent time series of multiple climate variables. They are among the most-used datasets in the atmospheric sciences and provide a comprehensive description of the observed climate as it has evolved during recent decades, on 3D grids at sub-daily intervals. Initially available at a resolution of $0.25^\circ \times 0.25^\circ$, the data were regridded to $1^\circ \times 1^\circ$ using Climate Data Operator (CDO) Version 2.5.0 (Schulzweida (2023)). The software is a collection of many operators for standard processing of climate and forecast model data developed by the Max Planck Institute for Meteorology. The adoption of regridding was based on the premise that low-frequency, large-scale atmospheric variability can be effectively captured without requiring high-resolution data. This approach ensures that the essential patterns of circulation remain identifiable while reducing computational costs. To further enhance the distinction between different circulation types, we worked out Z_{500} anomalies by subtracting the climatological mean from the actual data. The resulting anomaly field serves as the primary input for the SOM, with the resulting patterns used as the basis for the following drought analysis.

Monthly precipitation and surface temperature data were downloaded from the Climate Research Unit (CRU), TS version 4.08 (Harris et al. (2020)), already provided at a resolution of $1^\circ \times 1^\circ$ resolution. The CRU dataset extends from 1901-01-01 to 2023-12-01, covering a longer period than ERA5. SPI and SPEI indices were derived from the CRU datasets. The Python package `climate_indices` computes the climatic water balance as the difference between precipitation and PET, where PET is computed using the Thornthwaite equation (Thornthwaite (1951)). Commands `indices.spi` and `indices.spei` were employed to generate monthly datasets at a resolution of $1^\circ \times 1^\circ$, with different timescales such as SPI-1 (SPEI-1), SPI-6 (SPEI-6), and SPI-12 (SPEI-12).

Finally, monthly timeseries for NAO, EA, SCAND, and EAWR indices were obtained from the NOAA Physical Sciences Laboratory (PSL) (<https://psl.noaa.gov/data/climateindices/list/>), which conducts weather, climate, and hydrologic research to advance the prediction of water availability and extremes. The dataset covers a period from 1950-01-01 to 2024-10-01 and is derived from the NCEP/NCAR reanalysis (Kalnay et al. (1996)).

2.2 Technical details of the SOM algorithm

A SOM consists of a two-dimensional grid of nodes, where each node is associated to an N-dimensional weight vector, also known as a reference vector \mathbf{m}_i . To every \mathbf{m}_i corresponds to a two-dimensional pattern. The similarity between a data sample \mathbf{z} (a vector from the input dataset) and each reference vector is measured in terms of Euclidean distance. During training, the node with the smallest distance (eq 2.1), known as the Best Matching Unit (**BMU**), is identified as the closest match to the data sample:

$$\|\mathbf{z} - \mathbf{m}_{BMU}\| = \min\|\mathbf{z} - \mathbf{m}_i\| \quad (2.1)$$

This node is considered the "winner" for that iteration, effectively marking which pattern the input most resembles. Subsequently, only the vectors for the best-matching node and those that are topologically close to it in the two-dimensional array are updated by the relation

$$\mathbf{m}_i(t+1) = \mathbf{m}_i + h_{BMU,i}(t)[\mathbf{z}(t) - \mathbf{m}_i(t)] \quad (2.2)$$


The neighboring nodes are determined based on their proximity to the BMU on the grid, which can be defined by a neighborhood radius \mathbf{h} . The degree of the update is modulated by a neighborhood function. The most common neighborhood function is of Gaussian form (Skific and Francis, 2012), which

can be described as

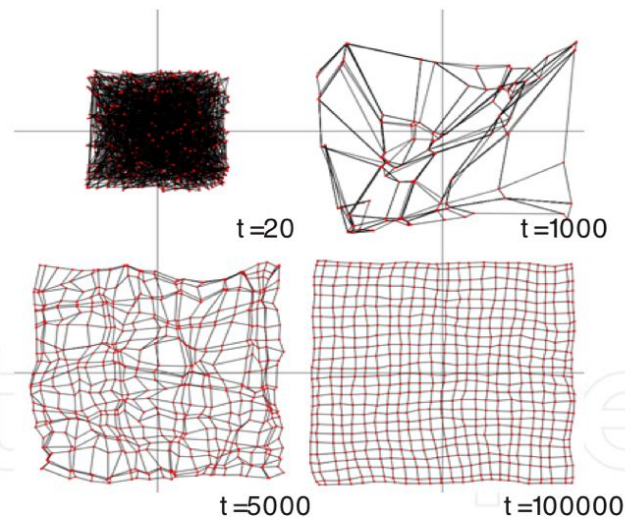
$$h_{BMU,i} = \alpha(t)e^{-\|r_{BMU}-r_i\|^2/2\sigma^2(t)} \quad (2.3)$$

where $\alpha(t)$ is the learning rate, \mathbf{r}_{BMU} and $\mathbf{r}_i \in \mathbf{R}^2$ are the radius vectors of the BMU and the i -th node respectively, and $\sigma(t)$ defines the width of the function (Johnson et al. (2008)). As training progresses, the quantization error \mathbf{q} is calculated to evaluate how well the SOM represents the input data. This metric measures the average Euclidean distance between each input vector and its corresponding BMU. Additionally, the topological error \mathbf{t} assesses how well the SOM preserves the spatial relationships in the input data. This error reflects the proportion of input samples for which neighboring BMUs are not located next to each other on the grid. A low topological error indicates that similar input samples are represented by adjacent nodes on the map, maintaining the structure of the original data distribution. Training continues iteratively until both the quantization error and the topological error reach acceptable levels. This process is often called convergence phase. Once the SOM is trained, the reference vectors of the nodes are mapped onto the two-dimensional grid (Fig 2.1). The patterns emerging on this grid reveal predominant states or clusters in the dataset, with each node representing a cluster of similar data points. Clusters located near each other in the grid correspond to similar patterns, while more dissimilar clusters typically appear at the corners of the map.

Once anomaly patterns are assigned to nodes, their frequency of occurrence (**FO**) can be determined, indicating how often specific patterns occur within the data. The frequency is a count of how many times each node has been the BMU, meaning it tracks how often a particular pattern is the closest match for any input vector. Nodes with higher activation frequencies correspond to patterns that are more common in the data, while nodes with lower activation frequencies represent less common patterns. Maps showing the FO for every node, as Fig 2.2 does, are called activation maps.



(b)



(c)

Figure 2.1: (b): SOM representation after the training. (c): Projection of the reference vectors for various stages of map training. From [Skific and Francis \(2012\)](#).

A critical aspect of SOMs is the determination of the number of nodes in the grid. This choice significantly influences the results, both in terms of pattern representation and in quantization error values and topographic error values, yet it remains inherently arbitrary, lacking a universally accepted guideline. Unlike other machine learning methods with clear criteria for model selection, SOMs heavily rely on domain expertise and the specific purpose of the analysis. A small number of nodes may be preferred for tasks like data compression or visualization, while a larger number of nodes might be necessary for clustering (Rousi et al. (2015)) or for identifying subtle patterns in complex datasets or for attribution studies (Odoulami et al. (2023)). However, this choice involves a delicate trade-off. A grid that is too small may lead to underfitting, where distinct clusters merge and important patterns are lost. On the other hand, an excessively large grid can result in overfitting, introducing artificial clusters that do not reflect meaningful structures in the data. Despite the existence of some "rules of thumb" which provides a starting point, researchers often rely on iterative testing, qualitative visualization, and metrics like quantization and topographic errors to fine-tune the grid size.

A commonly adopted guideline in attribution studies suggests selecting a sufficient number of nodes to ensure a continuous and well-distributed representation of atmospheric states, thereby preserving the structure of the data without introducing unnecessary fragmentation. Although these general guidelines exist, their application often leads to different choices in grid configuration, depending on the specific dataset and research objectives. For instance, while Harrington et al. (2016) opted for a relatively compact 3×4 grid (Fig. 1.6), Odoulami et al. (2020) and later Odoulami et al. (2023) employed a larger 6×5 grid (Fig. 1.7), whereas Rousi et al. (2015) chose an intermediate 4×5 grid. This variation highlights how the same foundational principles can lead to differing grid sizes, as the optimal number of nodes depends not only on theoretical considerations but also on practical factors

such as data complexity or spatial domain.




Figure 2.2: Activation map for the 7×5 SOM employed in this study. Each rectangle represents a node’s frequency of activation. As expected, the most frequent patterns are those in the corner of the map.

2.3 Methodology

Starting from the Python package [Minisom](#), a 7×5 SOM was developed. The adoption of a 7×5 grid is driven by the necessity for a continuous representation across an extensive spatial domain. Other configurations, including 3×3 , 4×3 , and 5×4 , did not achieve the desired optimal results. The implementation of a 7×5 map enabled the inclusion of a wider array of atmospheric states, many of which emerged from the superimposition of various atmospheric modes. This study specifically aimed to investigate also the states that correspond to overlapping configurations, as other approaches do not provide such insights. After choosing the number of nodes, the SOM must be trained and later optimized to minimize the quantization error over a subset of the entire Z_{500} dataset (70%), ensuring the best possible perfor-

mance. Lastly, the model is retrained using the whole dataset.

2.3.1 Preprocessing of Z_{500}

The preprocessing step is crucial for improving data quality by cleaning and normalizing it. In this specific case, monthly Z_{500} anomalies were first divided into a training set (70%) and a testing set (30%). Each subset was then standardized by dividing its values by their respective standard deviation, a process commonly referred to as normalization. Following normalization, Principal Component Analysis (PCA), a dimensionality reduction technique, was applied to the training set. PCA is a widely used method in machine learning and climate science to reduce the complexity of large datasets by transforming them into a smaller number of components that still capture the most significant patterns and trends. This step is fundamental as it simplifies the dataset, reducing noise and the risk of overfitting, especially when working with high-dimensional data such as large-domain atmospheric data. The use of PCA as a preprocessing step becomes particularly crucial when dealing with datasets that have more features than samples, a scenario commonly referred to as the ‘curse of dimensionality’ (Bellman (1961)). In such cases, PCA not only helps mitigate the curse of dimensionality but also enhances the data by focusing on the principal components that capture the majority of the variance. This focus enables noise reduction, leading to cleaner data and, potentially, more robust clustering results.

In this study, by using the Python package `sklearn.decomposition` and retaining 99% of the explained variance, the PCA identified 35 components. When applying PCA, the information is encoded into a reduced set, of components: at each time step, the dataset still corresponds to a unique spatial pattern, but instead of being represented in terms of raw Z_{500} values, it is represented in terms of a reduced set of principal component scores. When the SOM clusters these PCA-reduced representations, it is still effectively clustering spatial structures, just in a more efficient and compact form.

2.3.2 Training of the 7×5 SOM

The training step involves iteratively adjusting the network using the training set to minimize the quantization error. In this case, ten SOMs were generated, and after training, the SOM with the lowest q error was declared as the best performer.

Initially, the model was trained using arbitrary parameters, which provided a baseline performance. These parameters included settings such as the learning rate, the neighborhood radius, and the number of iterations, which influence how the SOM adapts to the data. It is important to underline that the first step of the training is to randomly initialize the SOM weights (e.g. the reference vector). In fact, it has been proven that SOMs results are not sensitive to the selected initialization method (Hewitson and Crane (2002)). Even after initialization, the training vectors keep being extracted randomly, to avoid falling into local minima. This choice reflects the no a priori assumption on the distribution of the data.

Following this initial training, a tuning process was conducted to optimize these parameters systematically. Tuning refers to the process of refining a model's hyperparameters to achieve better performance. This is typically done by exploring different combinations of parameter values. By fine-tuning the hyperparameters, the model becomes better equipped to generalize to unseen data while minimizing errors. In this case, the tuning process was critical to ensure that the selected SOM provided the most accurate representation of the input data with the lowest quantization error. The optimal set of parameters for the lower quantization error ($q = 66.37$) were:

1. Neighborhood radius $\sigma = 3$
2. Learning rate $lr = 0.001$
3. Number of iterations $n_{iter} = 10000$

Starting with the neighborhood radius, there is a "rule of thumb" that defines the right value for the neighborhood radius as $Y - 1$ (Hewitson and Crane (2002)). In the training step, σ was first set to $\sigma = 4$ and later changed to $\sigma = 3$ after the optimization step, as it returned a lower q error. The neighborhood function modulating the σ is the Gaussian function, because, as shown in Fig. 2.3, it was the function that minimized the q error. Its characteristic exponential decay ensures that nodes farther away from the Best Matching Unit (BMU), in terms of Euclidean distance, have progressively less influence.

The learning rate determines the step size taken during each iteration of the training process. It represents the speed at which a machine learning model "learns". If lr is set too high, the model may fail to converge, potentially leading to erratic behavior. Conversely, if the learning rate is too low, the model may converge extremely slowly or become stuck on incorrect configuration, effectively hindering learning. Typically, the learning rate is chosen within a range of values depending on the algorithm and the specific problem being addressed. For this model, values ranged from 0.001 to 0.0001.

Lastly, the number of iterations simply defines how many times the model has to be trained.

The full training procedure was reiterated with the comprehensive timeseries of data, but, this time, only the best performer was retrained. The final quantization error and topographic error values are

$$q = 66.49 \tag{2.4}$$

$$t = 0.01 \tag{2.5}$$

and the resulting activation map is shown in Fig 2.2. The following chapter will provide a detailed description of the remaining aspects.

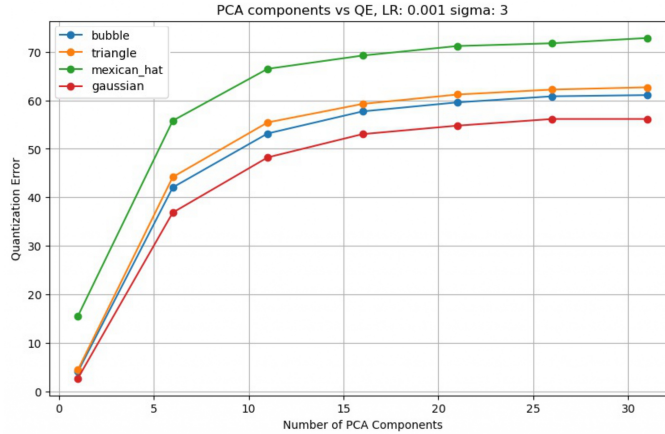


Figure 2.3: Q error evolution for different neighborhood functions.

2.4 Composite analysis and statistical significance testing

By analyzing the entire time period, the dates on which each node is the winner (i.e., the BMU) can be extracted. From this list of dates, the monthly and yearly occurrences can be shown. These kinds of characterizations allow for the identification of seasonal patterns that may be prominent in specific seasons or, on a yearly scale, distinguish between patterns that have been emerging or fading in recent years. This form of analysis could be integrated with results from other research examining drought conditions over extended periods, as it may help focus on recurrent patterns which could have connections to blocking states of the atmosphere.

To achieve a better characterization of each node, composite analysis was applied to calculate the composite mean values of the four teleconnection indices, precipitation, temperature, SPI-1, and SPEI-1. This statistical method

is a useful technique used in climate science to determine the basic structural characteristics of a meteorological phenomenon associated with a subset of the whole data. Composite analysis involves collecting large numbers of cases of a given phenomenon, which are then composited together as a collection. The composite analysis then generally involves computing the composite mean and perhaps computing some other statistical measures, such as the standard deviation and statistical significance.

By selecting only the dates in which the node activates and then computing the mean, this approach allows for the identification of a particular pattern, or more likely a specific region within the SOM, referred to as a cluster, which corresponds to a phase of a specific mode. First, the analysis was carried out considering the whole dataset as a continuous timeseries, then it was repeated dividing the year into extended winter (Dec-May) and extended summer (Jun-Nov). Further validation for the outcome was given by a statistical significance assessment, with the null hypothesis having a value significantly different from zero. Given the uniqueness of the dataset, characterized by the presence of a single timeseries, a bootstrap method was employed to assess statistical significance.

To assess the reliability of our results, we test statistical significance using bootstrap techniques. In research, statistical significance measures the probability of the null hypothesis being true compared to the acceptable level of uncertainty regarding the true answer. In particular, it helps determine whether the null hypothesis should be rejected or not. This analysis is widely used in research as it provides a quantitative measure of reliability and confidence. In this case, the null hypothesis is stated as follows:

The mean value is significantly different from 0.

To evaluate this hypothesis, it is essential to establish a distribution of mean values. Since the composite analysis yielded only one value, it became necessary to generate a distribution. A bootstrap technique was employed for

this purpose. Specifically, this method involves randomly selecting a value from the dataset, retaining it, and then reinserting it before drawing the next value. This cycle is repeated until a new dataset of the same size as the original is created. By constructing multiple resampled datasets, the bootstrap method approximates the sampling distribution of the statistic, allowing for higher confidence levels. Such techniques are particularly useful in cases where data is limited, as they provide a way to infer population parameters without requiring additional observations.

For this analysis, a total of 10000 bootstrap samples were generated, and the mean was calculated for each resampled dataset. A confidence interval was then determined by identifying the range between the 2.5th and 97.5th percentiles of the bootstrap distribution, corresponding to a 95% confidence level. Any observed mean falling outside this interval was considered statistically significant, suggesting that the original result is unlikely to have occurred by chance. Whenever this was the case, the mean value would be shown in bold font.

2.5 Analysis of drought events

Here we outline the methodology used to analyze significant drought events. As first step, we define the beginning and the end dates of each event. The time periods associated with these droughts vary in the literature, depending on the criteria used for drought classification. To ensure a systematic and objective definition, we adopted the following metrics:

- **Beginning of drought:** the month in which the SPI-12, averaged over the spatial domain of the drought, first reaches a value of 0.
- **End of drought:** the month corresponding to six months after the minimum value of the averaged SPI-12 within the same spatial domain.

This approach is justified by the fact that SPI-12 is calculated over a one-year

period, making it less sensitive to short-term fluctuations and seasonal variations. As a result, it provides a more stable and climatologically meaningful representation of drought onset and recovery. Additionally, the six-month period following the lowest SPI-12 value allows us to capture the lagged effects of drought persistence and hydrological recovery, which are critical for understanding long-term drought impacts.

After a brief overview of the surface variables during the event, we shift our focus to a more quantitative analysis of the relationship between the atmospheric patterns described by the SOM resulting patterns and precipitation. This characterization is done in terms of SPI-1 and SPEI-1 averaged over the reference spatial domain. We use SPI-1 in place of monthly precipitation anomalies as it provides an already standardized quantity easier to compare across different locations. Systematic comparison of SPI and SPEI should highlight the role played by high surface temperature in each drought event. We perform composite analysis for each node to find the distribution of the corresponding spatially averaged SPI-1 and SPEI-1 values, also computing the composite mean. To classify each node as, dry, wet or neutral, we employ a bootstrap technique which enables us to assess whether the composite mean SPI-1 and SPEI-1 values for a given node are significantly different from zero. Based on this analysis, we categorize the nodes into three groups:

- A node is classified as *dry* if the mean SPI-1 (SPEI-1) is significantly lower than 0.
- A node is classified as *wet* if the mean SPI-1 (SPEI-1) is significantly higher than 0.
- A node is classified as *neutral* if the mean SPI-1 (SPEI-1) is not significantly different from 0.

Once the nodes have been assigned to one of these three categories, we compute the frequency of occurrence for each class. This is done by counting

the number of months in which a dry, wet, or neutral node is activated during the event and then expressing these occurrences as a percentage of the total time period analyzed. This frequency provides information on the atmospheric configurations that prevailed during the drought.

To assess how exceptional the event was, we conducted the following analysis. Given the duration of the drought, we examined the entire dataset and identified all periods of equal length, same start, and same end month (for example, in the case of the Iberian drought, each period starting in November and ending in February of the following year). For each of these periods, we applied the same classification procedure described above to determine the activation frequencies of dry, wet, and neutral nodes. This process resulted in frequency distributions for each category, representing the typical variability across the dataset. We then compared these distributions with the activation frequencies observed during the drought event to determine how much they deviated from the expected range. Specifically, we calculated the percentile rank of the event's frequency within the distribution and derived the corresponding probability, quantifying how rare or unusual the observed frequencies were in a historical context.

Chapter 3

Results

In this chapter, we present the results of our analysis, which is divided into two main parts. First, we examine the Z_{500} patterns identified through the SOM algorithm, providing insight into the dominant atmospheric circulation regimes. We then characterize each node in terms of low-frequency variability and drought-related quantities, highlighting the connections between atmospheric variability and surface conditions.

In the second part, we apply the methodology introduced in Section [2.5](#) to three case studies (the 1974-1975 drought, the 2004-2005 drought, and the 2022-2023 drought) and analyze the atmospheric drivers of persistent droughts associated with each drought event. We highlight the SOM-based methods for drought analysis but also describe a framework and baseline for future attribution studies.

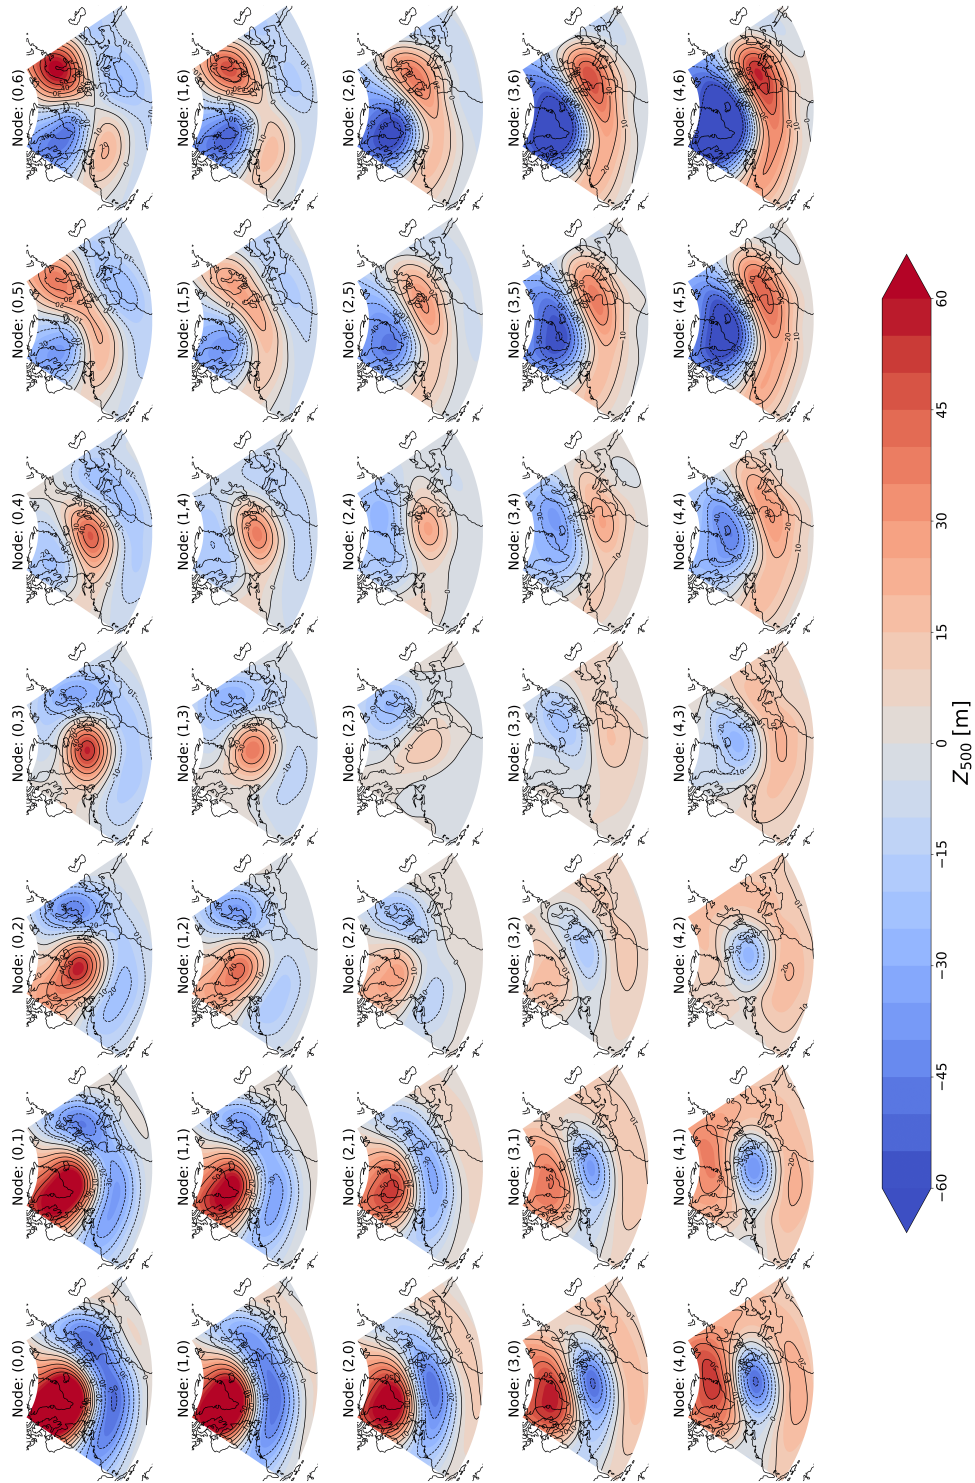


Figure 3.1: ERA5 Z_{500} monthly anomalies for 1940-2023 classified using the self-organizing maps (SOM) with 35 nodes.

3.1 SOM classification of atmospheric variability

Fig. 3.1 shows the 35 different patterns of monthly Z_{500} anomalies. The position of very distinct configurations at the corners of the SOM array demonstrates that the model has correctly placed the most different configurations as far apart as possible. Additionally, the activation map shown in Fig. 2.2 confirms that the model has correctly identified the most important patterns (e.g. those that occur most frequently) placing them at the corners of the map. The intermediate nodes display a gradual transition between extreme configurations, suggesting that the SOM effectively captures the continuum of atmospheric variability. Furthermore, the intensity of anomalies in the corner nodes highlights the presence of the most pronounced and contrasting patterns, reinforcing the validity of the classification. These intermediate configurations are crucial for this study, as they provide a comprehensive view of the variability that would be difficult to obtain with traditional methods such as EOF analysis.

The monthly and yearly occurrences for each node, are shown in Fig. A.3. These distributions highlight how certain patterns are predominantly associated with specific seasons, while others influence the synoptic scale uniformly throughout the year. Additionally, the variability in yearly occurrences (Fig. A.4) provides insight into possible long-term trends such as decadal variability or shifts in the frequency of specific patterns over time.

3.1.1 Relationship of SOM nodes to main teleconnections

In order to relate the SOM nodes to four main low-frequency teleconnections over the Atlantic-European sector, we work out the composite mean of the NAO, EA, SCAND, EAWR indices for each node (Fig. 3.2).

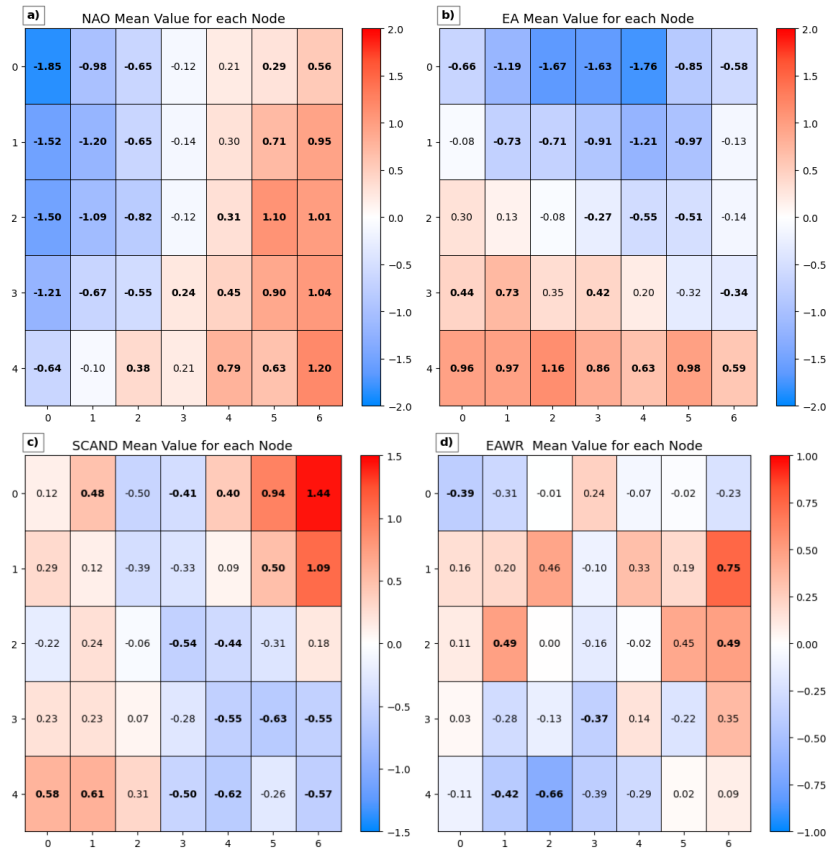


Figure 3.2: Mean value of the a) NAO, b) EA, c) SCAND, d) EAWR index for each SOM node. Bold values denote statistically significant from 0.

Mean values for the NAO index (Fig. 3.2a), show an area characterized by negative values (blue) in the upper left corner, and one characterized by positive values (red) in the opposite corner (bottom right). Most of the values are also significantly different from 0, although the magnitude of the positive and negative phases is not symmetric: the strongest positive value (1.20) is found in node (4,6), whereas the strongest negative value (-1.85) is in node (0,0). This asymmetry suggests that the NAO+ and NAO- phases may not have equivalent atmospheric impacts. These nodes exhibit patterns that closely resemble the canonical NAO structure (1.4), while nonsignificant

or close to zero values are associated with other atmospheric modes or a superimposition of them. Considering the monthly occurrences, Fig. 3.3 underlines that the NAO is most frequent during winter. This is true in both phases, although only one monthly frequency is shown.




Figure 3.3: Monthly frequency of the pattern corresponding to the node (0,0).

This is consistent with what described in scientific literature, that is the NAO is responsible for much of the low-frequency atmospheric variability in the North Atlantic region during the months of November to April (Barnston and Livezey (1987)).

The results for the extended summer and extended winter are shown in Fig. A.5. During winter, the NAO- area appears slightly reduced, covering fewer nodes compared to the annual analysis. The same cannot be said for the NAO+ area, which is more spread out in winter but appears more clearly defined during extended summer.

The mean values of the EA index (Fig. 3.2.b), are also grouped together in two very distinct areas corresponding to EA+ (bottom left), and EA- (upper center). With most of the nodes being statistically significant, the two phases are on opposite parts of the map. The node with the highest statistically significant value is node (4,2), while the most negative value is found in node (0,4). The negative phase shows overall stronger values than the positive phase, reaching almost -1.8, while the EA+, reaches 1.16.

Both patterns correspond to the typical structure for the EA (Fig. 1.4). Interestingly enough, when considering the yearly manifestation for the EA+ node, shown in Fig. 3.4, the frequency of occurrence of the pattern has increased since the beginning of the 21st century. This result is also confirmed by the standardized EA index, (see, e.g. https://www.cpc.ncep.noaa.gov/data/teledoc/ea_ts.shtml).




Figure 3.4: Yearly frequency of the pattern corresponding to the node (4,2).

Seasonal differentiation (Fig. A.6) shows a similar classification to Fig. 3.2.b during a) winter, while for b) summer we find less distinct results. In fact, node (1,6) presents a positive mean value even if belongs to the EA- area, although the value is not significant. EA- is stronger during winter, whereas EA+ is stronger in summer.

As far as the SCAND is concerned (Fig. 3.2.c), areas with well-defined values are less clear than in the case of NAO or EA. Although there is a distinct area for SCAND+ in the upper right corner and one for SCAND- in the lower right corner of this SOM array, significantly positive values can also be found in the leftmost part of the map. The positive phase is clearly stronger than the negative one: node (0,6) presents a mean value of 1.44, while node (3,5) has a mean of -0.63. This distribution suggests that the SCAND pattern in isolation appears less in the data with respect to the NAO or EA. Node (0,6) shows a geopotential pattern that clearly resembles the typical structure for the SCAND (Fig. 1.4), while for the negative phase, we struggle to find

correspondence.

Unlike NAO and EA, which are characterized by a dipole structure, the SCAND pattern features a more complex three-pole configuration, with centers over Scandinavia, the North Atlantic, and Eastern Europe (Fig. 1.4). This inherent complexity, along with its greater spatial and temporal variability, may explain why the SOM struggles to clearly isolate SCAND-related patterns from the data. Additionally, the presence of multiple interacting centers might lead to a more diffuse representation in the SOM, as reflected in the scattered distribution of near-neutral nodes and the weaker phase transitions across the map.

Seasonal analysis (Fig. A.7) further highlights this complexity. In winter, the SCAND- cluster spreads toward the center of the map, blurring the distinction between phases. Conversely, summer appears to offer the clearest separation between SCAND+ and SCAND- clusters, though substantial positive values can still be identified outside the upper right corner.

The SOM array for the EAWR (Fig. 3.2.d) shows even less distinct areas compared to the other teleconnections. Few nodes display significant values, and both positive and negative phases appear scattered across the SOM array without forming well-defined regions. A small concentration of EAWR+ can be observed in nodes (1,6), (2,5), (2,6), and (3,6), although only two of these nodes reach statistical significance. EAWR- appears primarily confined to the lower part of the SOM array, spanning from node (4,0) to node (4,4), with some additional presence in row 3. However, for both phases, isolated significant values also emerge outside these tentative clusters, further blurring the separation between positive, negative, and neutral states.

This blurring may also be partially explained by the intrinsic structure of the EAWR pattern, which is often described as a quadrupole pattern featuring alternating geopotential anomalies over the East Atlantic, western Europe, eastern Europe, and Western Russia (Fig. 1.4). This more complex and

spatially distributed structure could rarely appear in the data, leading to struggles in isolating EAWR-related patterns as clearly as it does for simpler, dipolar modes. These results also align closely with the findings from the EOF analysis. The EAWR pattern is associated with the fourth mode of variability, explaining a considerably smaller fraction of the total variance compared to the leading patterns, such as NAO and EA. This reduced importance is reflected in the SOM, as lower-variance EOF modes correspond to less dominant large-scale patterns, often characterized by weaker and more localized spatial signals. Thus, the imprint of EAWR on the atmospheric states is less robust, less persistent, and harder to capture for the SOM. Seasonal analysis (Fig. [A.8](#)) does not substantially clarify the clustering structure. A modest improvement is visible in summer, where EAWR- nodes are more consistently grouped in the bottom left corner of the map. However, EAWR+ remains diffuse across seasons, reflecting the overall weak phase separation.

We further estimate the probability of having a positive or negative NAO, EA, SCAND, and EAWR for each SOM node. A strong positive (negative) event is defined by a mean value exceeding 1 (-1). In practice, this is equivalent to considering ± 1 standard deviation, as the indices are standardized. The probability was computed simply by considering the number of successes over the total number of events.

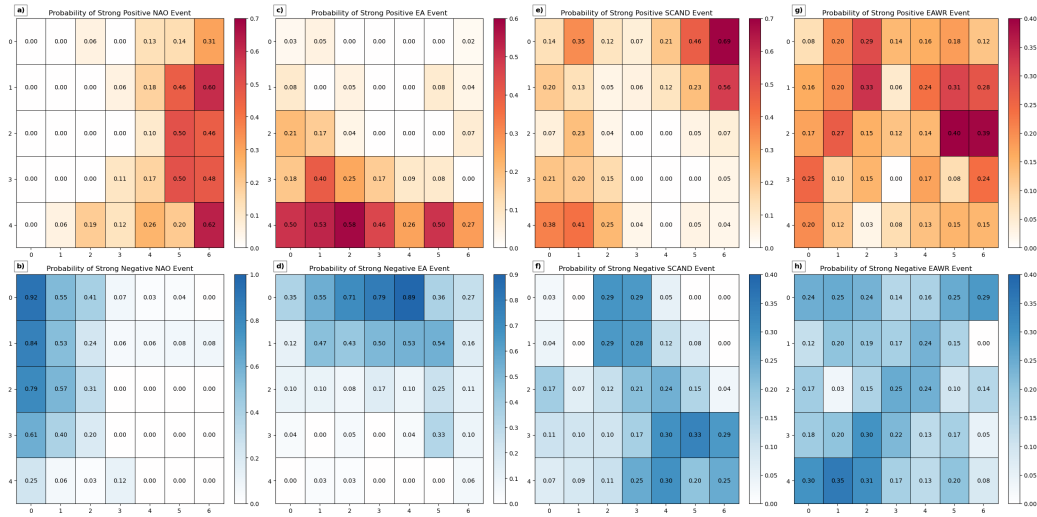


Figure 3.5: Probability array for a) NAO+, b) NAO-, c) EA+, d) EA-, e) SCAND+, f) SCAND-, g) EAWR+, h) EAWR- .

Fig. (3.5) shows a general consistency with the mean index value maps discussed in the previous subsections. Nodes characterized by positive mean values tend to exhibit a high probability of strong positive events, while nodes with negative mean values display a high probability of strong negative events. This relationship is particularly evident in the NAO (a and b) and EA (c and d) arrays. In these cases, the nodes with the highest probability of strong positive (negative) events coincide with the nodes associated with the most positive (negative) mean values. However, this consistency is partially lost in the SCAND array (e and f). The alignment holds for positive events, but the spatial coherence is reduced for strong negative events. The correspondence is almost completely disrupted in the EAWR (g and h) array, where there is little correspondence between mean values and event probabilities. This result was expected: as boundaries become less sharp in the mean value analysis, the corresponding probability maps also become more diffuse. The increasing difficulty in isolating coherent patterns for SCAND and EAWR reflects the weaker imprint of these modes on the atmospheric

variability, as already highlighted by the EOF analysis and the mean value distributions.

3.2 Relation between atmospheric patterns and drought-related quantities

To understand the influence and the impact of each individual node pattern over Europe, composite analysis of near-surface temperature precipitation anomalies, SPI-1, and SPEI-1 were evaluated. Results are shown for here temperature (Fig. 3.6), SPI-1 (Fig. 3.7), and SPEI-1 (Fig. 3.8), while precipitation maps can be found in Appendix A (Fig. A.9).

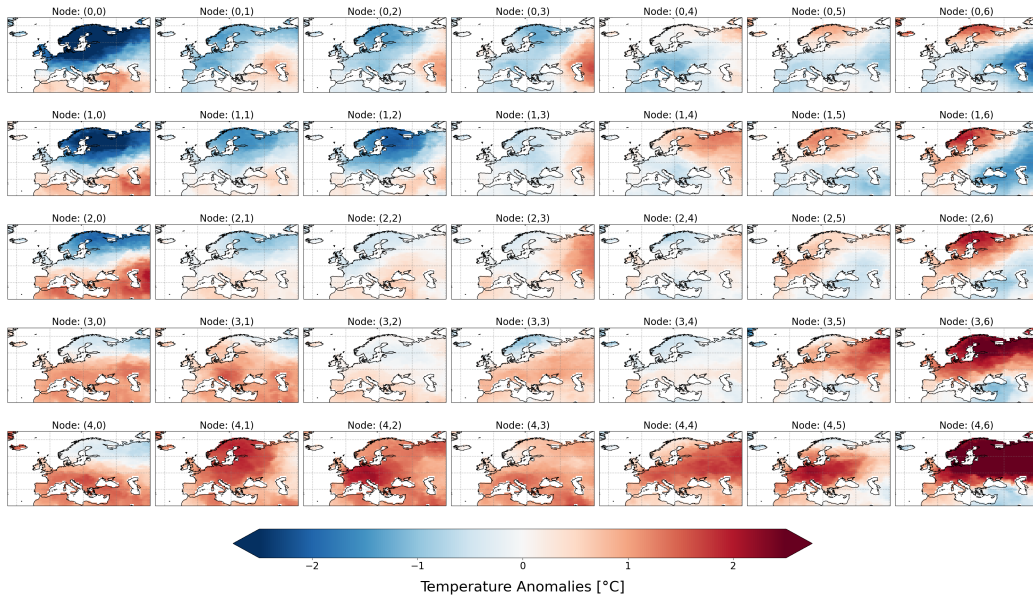


Figure 3.6: Composite of near-surface temperature anomalies for each SOM node.

In addition to composites, we also work out, for each SOM node, the prob-

ability of being in a dry month (i.e. SPI-1 or $\text{SPEI-1} \leq 1$). Specifically, the probability was computed simply as number of successes (i.e. having a value under the threshold) over the total number of events. The resulting maps are shown in Appendix A (Fig. [A.10](#) and Fig. [A.11](#)).

The temperature anomaly maps clearly reflect the influence of the different geopotential patterns over Europe. Consistent with expectations and evidence illustrated in the Assessment Report 6 (AR6), WG1 ([IPCC \(2023\)](#)), the NAO+ group leads to widespread positive temperature anomalies especially over northern Europe, while NAO- nodes feature negative anomalies over the same region. A similar, though less pronounced, relationship is evident for the EA nodes. An interesting resulting pattern emerges from the SCAND+ pattern (upper right), where the positive temperature anomaly is concentrated over the Scandinavian Peninsula and negative temperature anomalies are evident over southern-eastern Europe.

Beyond the specific phases of individual atmospheric modes, a broader and more notable feature emerges: patterns indicative of stationary Rossby waves, characterized by persistent high-pressure systems over Europe, consistently correspond to positive temperature anomalies over large portions of the continent. Conversely, cyclonic and rain-bearing configurations are generally associated with negative temperature anomalies. This coherence between the geopotential fields and surface temperature response highlights the strong coupling between large-scale circulation patterns and near-surface climate variability in the Euro-Atlantic region.

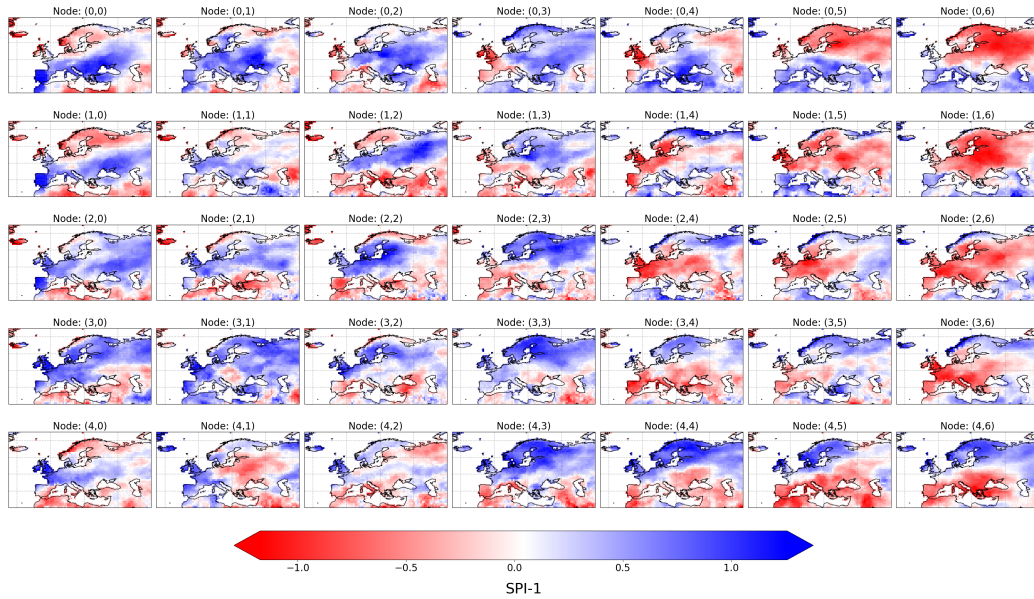


Figure 3.7: Composites of SPI-1 for each SOM node.

SPI-1 maps (Fig. 3.7) suggest that temperature alone cannot fully explain the observed precipitation anomalies. While there is a general correspondence between colder-than-average (warmer-than-average) conditions and drier (wetter) conditions—particularly in the corners associated with NAO- and NAO+ (nodes (0,0) and (4,6))—this relationship does not hold consistently across all patterns. For instance, when considering the opposite corners (nodes (0,6) and (4,0)), the link between temperature anomalies and precipitation variability appears weaker. This suggests that other factors play a crucial role in modulating drought conditions beyond what temperature alone can account for. Although the transition is less straightforward than in the temperature maps, the intermediate nodes play a crucial role in ensuring a smooth and gradual shift between the atmospheric patterns represented in the corner nodes, capturing the continuum of variability in precipitation responses.

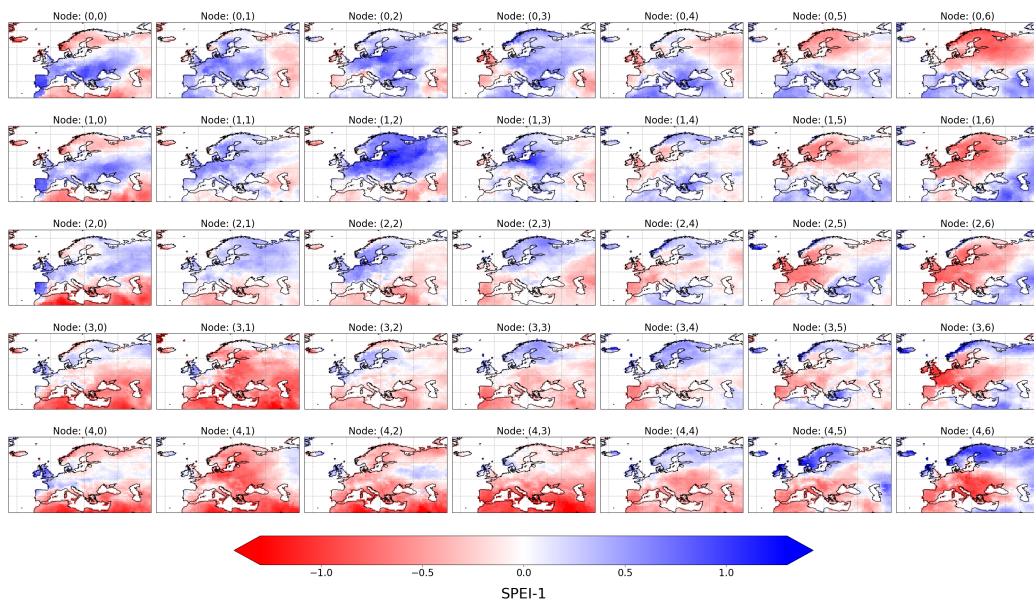


Figure 3.8: Composites of SPEI-1 for each SOM node.

When examining the SPEI-1 maps (Fig. 3.8), we gain a more comprehensive view of the relationship between atmospheric patterns and precipitation and temperature anomalies. In fact, the SPEI-1, thanks to the inclusion of evapotranspiration, provides additional information that helps explain responses that SPI-1 alone could not justify. This demonstrates that comparing SPI and SPEI can be highly valuable in studies of this kind, as it allows for a more complete understanding of the mechanisms driving hydroclimatic variability. For this reason, we chose to use both indices as criteria for the analysis conducted in Section 3.3.

A particularly noteworthy case is the EA+ cluster, located in the lower-left part of the SOM array. As previously illustrated in Fig. 3.4, this pattern has shown an increase in frequency since the 2000s. In this context, the EA+ cluster is associated with notably strong negative values for the SPEI-1 across large portions of southern and central Europe. However, as depicted

in Fig. 3.7, the same pattern in earlier years appeared to have a less pronounced impact on the composite. Under SPEI-1 it is possible to highlight the growing connection between this specific atmospheric configuration and the increased risk of drought in recent decades.

The consistency observed across the geopotential, temperature, and precipitation maps highlights the potential of SOMs as a valuable tool in drought attribution research. By effectively capturing the relationships between large-scale atmospheric patterns and surface climate anomalies, SOMs demonstrate their capacity to extract key information relevant to drought development and persistence. Their ability to identify both canonical and transitional atmospheric states, along with the associated surface responses, suggests that their application could offer new insights into the complex dynamics of droughts, revealing connections that may be overlooked by more traditional approaches.

3.3 SOM-based analysis of atmospheric drivers of drought events

This section presents findings related to three major, multi-year European drought events: the 1975–1976 drought in the United Kingdom, the 2004–2005 drought in the Iberian Peninsula, and the 2022–2023 drought in Southern Europe. To provide context, we first present a brief overview of each drought, outlining its severity, duration, and socioeconomic consequences. We then introduce the results obtained through the methodology described in Section 2.5, which offers insights that traditional approaches, such as simple correlations between atmospheric and surface variables or standard statistical tools, fail to capture. Through the implementation of a comprehensive SOM-based framework across the three events, this study gives a new insight into atmospheric drivers of these droughts and provides a general framework for subsequent attribution studies.

3.3.1 The 1975-1976 drought

The 1975-76 drought is considered one of the most severe experienced across much of the UK. Although shorter in duration relative to other major droughts experienced in the last century, the 1975–76 drought is widely acknowledged as the benchmark drought of the last 50 years in the UK, with no other historical drought matching the degree of spatial coherence, geographic extent or hydrological intensity (Parry et al. (2012)). It caused considerable depletion of surface water and groundwater over most of England and Wales as well as neighboring parts of the Continent, while Scotland and Northern Ireland were relatively less affected (Marsh et al. (2007)). The drought seriously reduced agricultural production, caused the industry to reuse water, restricted navigation, it damaged buildings, made householders become very cautious in their use of water, and provided a severe test of the water authorities (Rodda and Marsh (2011)).

This remarking event started in May 1975 and ended in late autumn 1976, with the hydrological drought lagging behind rainfall deficiencies over these 16 months (Parry et al. (2012)). It initially developed gradually over the winter months, primarily affecting the UK, especially the southeast, as well as, to a lesser extent, southern Scandinavia. Throughout the winter, the drought became increasingly spatially coherent across most affected regions. In the UK, the earliest impacts were felt in the western regions from October to December 1975 before the drought shifted eastward, where it remained concentrated until eventually spreading across the entire country and reaching the continent by the summer of 1976 (Parry et al. (2012)). By January 1976, the most severe deficits were observed in southeastern UK. The delayed onset in this region can be attributed to its higher groundwater storage capacity, which initially provided resilience against short-term precipitation deficits. Regarding large-scale atmospheric variability, the evolution of the NAO during the event does not adequately explain the precipitation deficit, as the period was marked by neutral to moderately positive values. In contrast,

the EAWR exhibited the most significant deviations from neutrality, though these shifts toward the positive phase remained moderate (Parry et al. (2012)). Although this drought has significantly impacted also other regions across Europe, here we focus on the British Isles as the area of interest and define the spatial domain as $[12^{\circ}\text{W}, 4^{\circ}\text{E}, 50^{\circ}\text{N}, 60^{\circ}\text{N}]$. Using the methodology described in Section 2.5, we set the beginning of the drought in August 1975 and its end in February 1977.

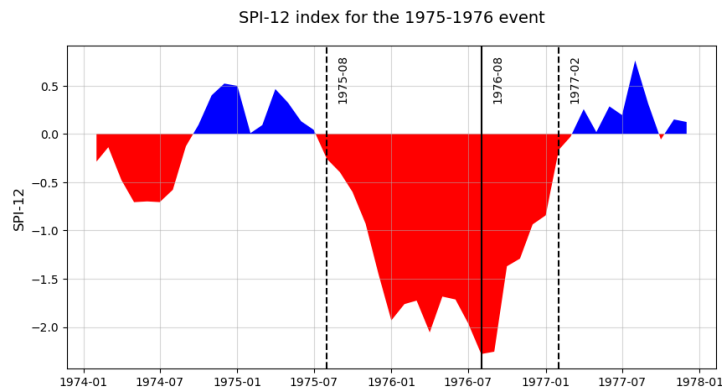


Figure 3.9: Spatially averaged SPI-12 over the UK domain. Solid line shows the UK domain, dashed lines correspond to the beginning and end dates for the event.

The spatial distribution of SPI-12 and SPEI-12 for August 1976, shown in Fig. 3.10, represents the conditions during the month that recorded the lowest value as depicted in Fig. 3.9. Both indices indicate critical drought conditions across the UK, with negative values dominating the maps. However, SPI-12 shows more extreme drought conditions, with larger areas experiencing values below -2.5, particularly in southern England. In contrast, SPEI-12 exhibits a less intense pattern, with fewer regions reaching the most severe drought classifications.

The fact that SPI-12 depicts dryer conditions suggests that the drought was primarily driven by a significant lack of precipitation rather than by an increased atmospheric evaporative demand, meaning that precipitation deficits

alone were enough to cause extreme drought conditions. However, although evapotranspiration does not appear to be the main driver of the drought, it may have played a role in modulating the severity in some regions, particularly in southern England, where the drought was more intense and where the lowest values of SPEI-12 are found. This suggests that in these areas, the increase in atmospheric demand could have contributed to the intensification of dry conditions.

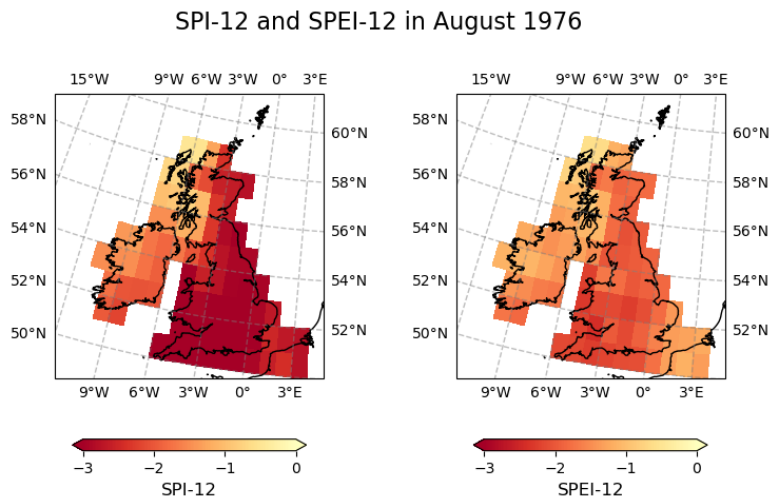


Figure 3.10: Spatial distribution of SPI-12 (left) and SPEI-12 (right) in August 1976, the month corresponding to the minimum value identified in Fig. 3.9.

When characterizing precipitation anomalies in terms of SOM nodes, several interesting patterns emerge. Using the entire time series and considering only the dates corresponding to the activation of each node, the occurrence of negative and positive precipitation anomalies was computed over the study area. The number of months in which the mean precipitation anomaly is negative for each node is shown in Fig. 3.11.a, while Fig. 3.11.b shows the number of occurrences for positive anomalies. Finally, Fig. 3.11.c shows the mean precipitation anomaly over the entire time period, providing a direct measure

of the average deviation from normal precipitation conditions. Notably, the nodes with the highest frequencies do not necessarily correspond to those associated with the most extreme positive or negative precipitation anomalies. This highlights a clear distinction between frequency and intensity: circulation patterns that occur frequently are not always linked to the most severe precipitation anomalies. A striking example is node (4,5), which, despite having far fewer positive anomaly occurrences than nodes (4,0) and (4,6), exhibits a mean composite value in Fig. 3.11.c of nearly the same magnitude. Another particularly interesting case is node (0,6), which displays a nearly equal number of negative and positive occurrences, resulting in a mean precipitation anomaly close to zero. However, the slightly higher number of positive occurrences suggests that negative anomalies tend to be stronger.




Figure 3.11: Occurrences of a) negative precipitation anomalies and b) positive precipitation anomalies over the UK study area and c) mean precipitation anomaly associated with each SOM node.

Aggregating nodes associated with precipitation anomalies of the same sign, we have a characterization of dry, wet, and neutral nodes for the UK in terms of a) SPI-1 and b) SPEI-1 (Fig. 3.12). The classifications provided by the two indices are largely similar and consistent with the patterns observed in Fig. 3.11.c, confirming the overall coherence in the spatial distribution of dry

and wet conditions. However, some differences emerge when comparing the two indices. SPI-1 identifies a total of 9 dry nodes and 11 wet nodes, while SPEI-1 detects 10 dry nodes but only 6 wet nodes. This discrepancy implies that several nodes classified as wet by SPI-1 are instead categorized as neutral when considering SPEI-1. This result suggests that in these nodes, evapotranspiration may have counterbalanced the precipitation excess, leading to a reclassification from wet to neutral.




Figure 3.12: Characterization of wet (blue), dry (red), and neutral (yellow) SOM nodes based on a) SPI-1 and b) SPEI-1 for the UK study area.

After characterizing the dry, wet, and neutral nodes in terms of SPI-1 and SPEI-1, we computed the activation frequencies of the three categories for the period going from August 1975 to February 1977. Table 3.1 indicates that, in terms of SPI-1, the dry nodes were active 57.89% of the time, the wet nodes 10.53%, and the neutral nodes 31.58%. Regarding SPEI-1, the dry nodes were active 57.89% of the time, the wet nodes 5.26%, and the neutral nodes 36.85%. Although the frequency of dry nodes remains unchanged, the activation of wet nodes is halved compared to SPI-1, leading to a corresponding increase in neutral nodes. This behavior is expected, as the additional sensitivity of SPEI-1 led to a lower number of nodes classified as wet. On the other hand, the number of dry nodes remains nearly the same (Fig. 3.12). Fig. 3.13 presents the activation frequencies of dry, wet, and neutral nodes

for a) SPI-1 and b) SPEI-1 during the 1975-1976 drought, compared against the distribution of values during 1940-2023 due to natural climate variability. Each boxplot represents the distribution of activation frequencies computed over all periods of the same start, same end, and equal length to the event, providing a benchmark for typical variability. The red dots indicate the specific activation frequencies observed during the 1975-1976 drought. For both metrics, the frequency of activation of dry nodes during the event is higher than the median of the reference distribution (which is 28.95% for both), indicating a greater-than-expected persistence of drought conditions. The wet category shows an even more striking deviation, with activation frequencies during the drought significantly lower than the reference values, reinforcing the notion of a pronounced and sustained precipitation deficit. Meanwhile, the neutral nodes also display deviations from typical values.




Figure 3.13: Frequency distribution for dry, wet, and neutral nodes considering all the periods with the same start (end) month. Red dots represent the observed frequencies for the event. The box spans from the 25th to the 75th percentile, with a line at the median. Whiskers extend from $25Q - 1.5 \times IQR$ to $75Q + 1.5 \times IQR$, where $IQR = 75Q - 25Q$.

Table [3.1](#) presents the observed activation frequencies, their corresponding

percentiles, and the associated probability. The percentile was calculated empirically, comparing the observed frequency with the reference distribution. Additionally, the probability was computed in terms of obtaining an equally or more extreme value, representing the likelihood of observing such an extreme activation frequency under typical conditions. The probability is equal to the percentile when considering the lower tail of the distribution; otherwise, it is given by $1 - \text{percentile}$ when referring to the upper tail.

	SPI-1			SPEI-1		
	Freq. (%)	Pctl	Prob (%)	Freq. (%)	Pctl	Prob (%)
Dry	57.89	98.36	1.64	57.89	99.65	0.35
Wet	10.53	0.75	0.75	5.26	0.01	0.01
Neutral	31.58	86.67	13.33	36.85	89.47	10.53

Table 3.1: Observed frequencies, percentiles, and probabilities for SPI and SPEI during the 1975-1976 drought.

The results from the probability values confirm that the 1975-1976 drought was indeed an exceptional event, both when considering SPI-1 and SPEI-1 indices. The corresponding percentiles for the dry categories are extremely high, 98.36 for SPI-1 and 99.65 for SPEI-1, reflecting a almost unprecedented persistence of dry conditions. Notably, the probability of encountering dry conditions is lower for SPEI-1 (0.35%) compared to SPI-1 (1.64%), indicating that the incorporation of evapotranspiration in the latter intensifies the severity of the drought. For the wet category, the probabilities are similarly low: 0.75% for SPI-1 and 0.01% for SPEI-1, further confirming the prolonged dryness during this period. The neutral category, with probabilities of 13.33% for SPI-1 and 10.53% for SPEI-1, suggests that these neutral atmospheric states, while not directly related to extreme dryness or wetness, still contributed to the overall atmospheric conditions.

In conclusion, these findings underscore the rarity and significance of the 1975-1976 drought, with probability estimates pointing to an outcome far

beyond traditional assessments. Thanks to this new approach, which has never been applied to this specific case, we have also quantified the distinct contributions of dynamical and thermodynamical drivers.

3.3.2 The 2004-2005 drought

The 2004-05 hydrological year was marked by one of the most severe droughts ever recorded in the Iberian Peninsula, particularly in the central and southern regions. Intense droughts are a common feature of the Iberian precipitation regime, which is characterized by strong interannual and decadal variability (Esteban-Parra et al. (1998)). However, the low levels of precipitation and the widespread impact have made this event one of the most devastating episodes in the region's history. According to Spanish and Portuguese authorities, several long-term precipitation stations across central Spain and Portugal recorded the worst drought since the late nineteenth century. In southern Iberia, it ranks as the second most intense drought in the last 130 years. This severity led to significant political and social unrest, as water levels in dams across Spain and Portugal dropped to around 55% of their total capacity by June 2005 (García-Herrera et al. (2007)). River flow values throughout Iberia also dropped substantially, reaching record lows, with direct consequences for hydroelectric power production. Iberian precipitation is highly variable across the region, and even the most widespread droughts do not affect the entire territory. Typically, about 20% of the Iberian peninsula is impacted during a drought event, with the most intense episodes affecting over 50% of the territory (Vicente-Serrano and López-Moreno (2008)).

The NAO plays a key role in controlling precipitation variability in the region, particularly during winter. At the same time, other large-scale atmospheric circulation patterns, such as the EA and SCAND patterns, also influence precipitation in Iberia. Although García-Herrera et al. (2007) demonstrated that scarce precipitation in the Iberian peninsula is closely linked to the positive phase of the NAO, its influence alone did not fully account for the low

precipitation levels observed in February. Instead, a negative phase of the EA provided a more convincing explanation. In contrast, March 2005, stands as an exception: despite the negative NAO and positive EA, neither index adequately explained the significant negative precipitation anomalies over much of Europe. The primary driver of this unusual precipitation pattern was a blocking episode in early March, which caused atypical shifts in storm tracks, preventing cyclonic activity from reaching Iberia. Such blocking events during negative NAO phases are less frequent over Iberia compared to other regions, as the displacement of blocking centers in this case favored dry conditions across the Peninsula. This highlights the importance of a SOM-based approach in this context, as it provides a more flexible method to identify atmospheric patterns driving droughts.

For this case study the spatial domain for averaging SPI and SPEI is defined by $[10^{\circ}\text{W}, 5^{\circ}\text{E}, 35^{\circ}\text{N}, 45^{\circ}\text{N}]$. Given the time evolution of the area-averaged SPI-12 (Fig. 3.14), we consider the period November 2004 - February 2006 as characteristic of the drought duration.

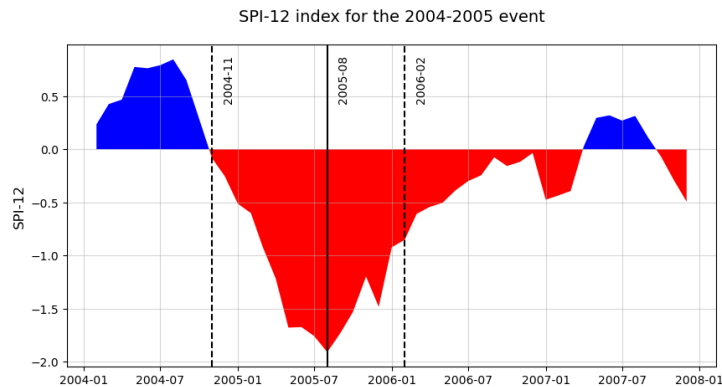


Figure 3.14: Spatially averaged SPI-12 over the Iberian domain. Solid line shows the month corresponding to the minimum value, dashed lines correspond to the beginning and end dates for the event.

The spatial distribution of SPI-12 and SPEI-12 is shown in Fig. 3.14 for the month of August 2005, that is the time of maximum dryness (Fig. 3.14).

Both indices indicate critical drought conditions across the Iberian Peninsula, however, SPI-12 exhibits more extreme drought conditions in several regions, particularly in the central-eastern parts of the peninsula. In contrast, SPEI-12 presents a slightly less severe pattern, with fewer areas reaching the most extreme drought classifications.

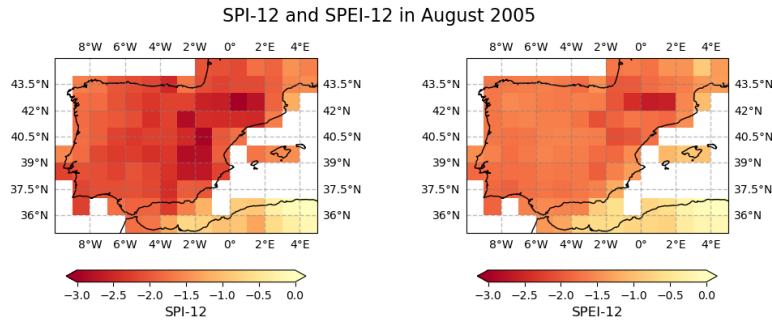


Figure 3.15: Spatial distribution of SPI-12 (left) and SPEI-12 (right) in August 2005, month corresponding to the minimum value identified in Fig. 3.14.

The characterization of precipitation anomalies in terms of the SOM nodes, shown in Fig. 3.16, reveals several interesting features. When considering occurrences of negative anomalies, a key difference from the previous case is that the highest frequency and highest intensity now coincide in the same node, specifically node (4,6). However, this is not the case for positive anomalies in Fig. 3.16.b, where node (0,6) does not correspond to the highest mean precipitation anomaly. Instead, other nodes, such as (4,0) and (4,5), exhibit greater positive anomaly magnitudes despite occurring less frequently. This suggests that while some circulation patterns lead to frequent wet conditions, they are not necessarily associated with the most extreme precipitation events.

Additionally, when comparing Fig. 3.16.a and Fig. 3.16.b, the spatial distribution of positive and negative anomalies appears asymmetric. Negative anomalies do not exhibit any clear clustering, whereas two major regions

emerge for positive anomalies: one in the upper right corner and another along the left side of the array. Finally, Fig. 3.16.c shows a stronger resemblance to Fig. 3.16.b than to Fig. 3.16.a, reinforcing the idea that positive precipitation anomalies tend to be more spatially structured, whereas negative anomalies are more evenly distributed.

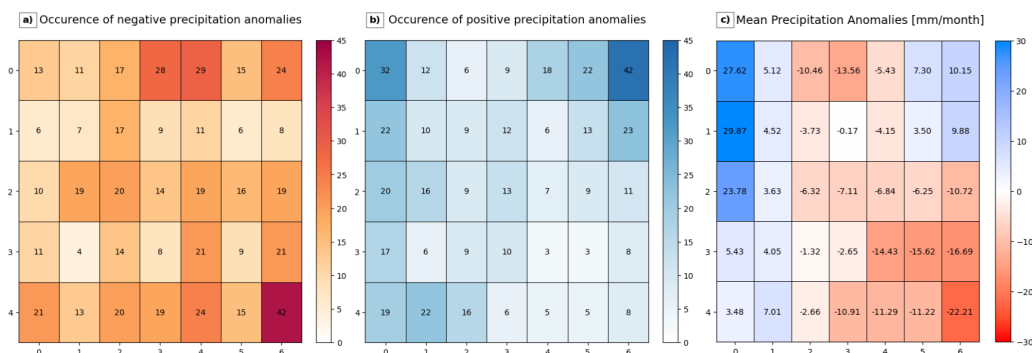


Figure 3.16: Occurrences of a) negative precipitation anomalies and b) positive precipitation anomalies over the Iberian peninsula study area, and c) mean precipitation anomaly associated with each SOM node.

The characterization of dry, wet, and neutral nodes in terms of a) SPI-1 and b) SPEI-1 is shown in Fig. 3.17. Compared to the previous case, the classifications here exhibit greater discrepancies between the two indices. While both maps show some resemblance to Fig. 3.16.c, the agreement between SPI-1 and SPEI-1 is less clear. In particular, the identification of wet nodes appears relatively consistent between the two and the precipitation anomalies, whereas the correspondence for dry nodes is less straightforward. SPI-1 identifies only 8 dry nodes and 7 wet nodes, whereas SPEI-1 detects a total of 16 dry nodes and just 5 wet nodes. This difference suggests that SPEI-1 classifies as dry even atmospheric patterns that are not associated with large negative precipitation anomalies, significantly increasing the number of dry nodes while reducing the neutral ones. The fact that SPEI-1 iden-

tifies nearly twice as many dry nodes as SPI-1 implies that in a southern European region such as Iberia, evapotranspiration plays a crucial role in intensifying drought conditions. Indeed, Iberia is naturally prone to hot, dry summers, where high temperatures and increased evapotranspiration exacerbate drought conditions. SPEI-1 better captures these effects related to soil drying, thus enlarging the number of dry nodes and reducing the number of neutral nodes.

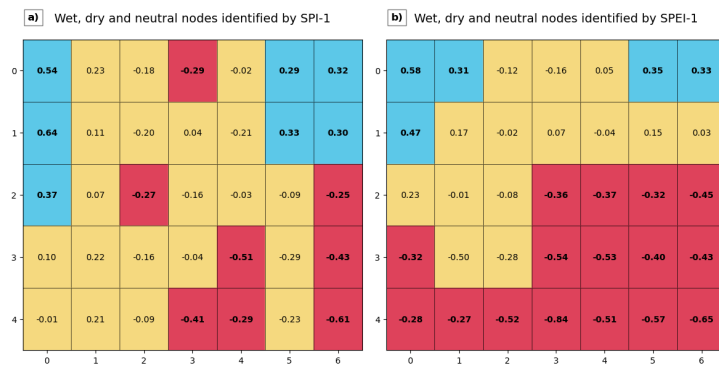


Figure 3.17: Characterization of wet (blue), dry (red), and neutral (yellow) nodes based on a) SPI-1 and b) SPEI-1 for the Iberia study area.

During the drought, the activation frequencies for each class were found to be 33.33% (dry), 16.67% (wet), and 50.00% (neutral) for SPI-1, while for SPEI-1, the frequencies were 41.67% (dry), 25.00% (wet), and 33.33% (neutral). Fig. 3.18 presents them in comparison to the distribution of values during 1931-2023 due to natural climate variability. Under SPI-1, the frequency of dry node activation during the event falls within the interquartile range (IQR), suggesting that dry conditions were broadly in line with typical variability when evaluated from the perspective of precipitation alone. However, this is not the case for wet nodes, where the activation frequency during the event lies far below the IQR, indicating a notably lower occurrence of wet conditions than expected.

Under SPEI-1, this relationship is inverted: the dry node frequency is more

extreme, aligning with the upper tail of the reference distribution, which suggests a pronounced dry anomaly when accounting for evapotranspiration effects. Conversely, wet node activation is below the IQR but remains within the whiskers, indicating a marked reduction in wet conditions, though not at an extreme level. Neutral node frequencies under SPEI-1 are closer to the typical range, suggesting a more balanced distribution of neutral conditions compared to SPI-1. The differences in frequency for dry and neutral nodes between SPI-1 and SPEI-1 reflect the difference in classification of the nodes, as many nodes that for SPI-1 were classified as neutral resulted dry for SPEI-1.




Figure 3.18: Frequency distribution in terms of a) SPI-1 and b) SPEI-1 for dry, wet, and neutral nodes considering all the periods with the same start (end) month. Red dots represent the observed frequencies for the event. The box spans from the 25th to the 75th percentile, with a line at the median. Whiskers extend from $25Q - 1.5 \times IQR$ to $75Q + 1.5 \times IQR$, where $IQR = 75Q - 25Q$.

Table 3.2 presents the observed activation frequencies, their corresponding percentiles, and the associated probabilities. Despite the event being recognized as a significant drought, the probability associated with the activation frequency of dry nodes under SPI-1 (32.43%) suggests that dry conditions were relatively typical, as it falls near the median of the reference distribution. However, SPI-1 effectively captures the anomalously low activation of wet nodes, with a probability of only 9.05%, highlighting its deviation from reference variability. What stands out in this case is the exceptional probability associated with neutral nodes (0.14%), indicating that these configurations played a dominant role in shaping the event. This suggests that the drought, as depicted by SPI-1, was not driven by an extreme increase in dry conditions, but rather by an unusually high presence of neutral conditions combined with a significant reduction in wet occurrences.

SPEI-1, on the other hand, places stronger emphasis on the dry anomaly, as reflected in the high percentile (93.75) and relatively low probability (6.25%), suggesting that the activation of dry nodes was much less typical compared to SPI-1. However, SPEI-1 is less effective in capturing the rarity of the low wet-node activation, which appears more anomalous in SPI-1 (9.05%) than in SPEI-1 (19.48%). The activation frequency of neutral nodes under SPEI-1 does not appear particularly unusual, with a higher probability (21.43%) compared to SPI-1. This suggests that while SPI-1 highlights the anomalous reduction in wet conditions, SPEI-1 better identifies the extent of the dry anomaly, likely due to its incorporation of evapotranspiration effects.

	SPI-1			SPEI-1		
	Freq. (%)	Pctl	Prob (%)	Freq. (%)	Pctl	Prob (%)
Dry	33.33	67.57	32.43	41.67	93.75	6.25
Wet	16.67	9.05	9.05	25.00	19.48	19.48
Neutral	50.00	99.86	0.14	33.33	78.57	21.43

Table 3.2: Observed frequencies, percentiles, and probabilities for SPI-1 and SPEI-1 during the 2004-2005 drought.

To conclude, the analysis of the 2004-2005 drought provides valuable insights through the SOM-based characterization. The results highlight a key difference between SPI-1 and SPEI-1 in how the drought was represented. Under SPI-1, the event was primarily driven by an unusually high activation of neutral nodes, rather than an extreme increase in dry conditions. In contrast, SPEI-1 attributes the drought more directly to the activation of dry nodes, emphasizing the role of temperature in enhancing evapotranspiration.

3.3.3 The 2022-2023 drought

In 2022, intense and prolonged drought conditions severely affected large portions of France, Italy, and Spain, with widespread socio-economic and ecological consequences. The drought was driven by a persistent precipitation deficit that began in late 2021, becoming particularly evident in northwestern Italy by March 2022. Poor snow accumulation in the Alps and unusually warm winter temperatures raised concerns about reduced snowmelt contributions to river flows in spring, increasing the risk of hydrological drought (GDO report for March 2022). Drought severity indicators such as SPI and SPEI frequently fell below -2 (categorized as extreme drought with respect to Tab. 1.1), with some areas recording values below -3 (Faranda et al. (2023)). The impacts were severe across sectors: in Italy, water restrictions affected around 50% of the population, especially in the North. Record inland saltwater intrusion was reported in the Po River Delta, reaching up to 40 km from the coast. Agriculture suffered significant losses, with maize, soybean, and sunflower yields reduced by approximately 15% compared to the 5-year average (GDO report for March 2022). Hydropower production and cooling systems for power plants were also strained due to reduced water availability. Similar impacts were observed in France and Spain. Southern France experienced extensive wildfires, with burned land more than double that of 2021 and 4.6 times the 2012–2021 average. By August, 66 French departments were under the highest drought alert level, with 93 under one

of the top three levels. In Spain, Portugal, and the Netherlands, agriculture, energy production, and domestic water supply were similarly affected. The exceptional nature of the 2022 drought drew considerable media attention, elevating water management as a priority on the political agenda. This event also fueled public debate on the role of anthropogenic climate change. Studies such as Faranda et al. (2023) and Bevacqua et al. (2024) suggest that ACC intensified the drought through a thermodynamic effect, making the atmosphere both ‘stronger’ and ‘warmer,’ thereby exacerbating drying conditions. However, no clear evidence was found of changes in the frequency of the slow-evolving atmospheric circulation pattern that triggered the drought, indicating that while global warming amplifies the intensity of drought, it has not significantly altered its patterns of occurrence. Although this exceptional drought affected large portions of central and western Europe, the spatial domain we chose for area-averaging is [7°W, 17°E, 35°N, 49°N] (Fig. 3.20), which corresponds to the part of Europe where the drought lasted longer and the lowest values of drought indices were reached. The time period identified by the spatially averaged SPI-12 spans from March 2021 to February 2023, with a total duration of 24 months.




Figure 3.19: Spatially averaged SPI-12 over southern Europe domain. Solid line shows the month corresponding to the minimum value, dashed lines correspond to the beginning and end dates.

The spatial distribution of SPI-12 and SPEI-12 for August 2022 (Fig. 3.20) illustrates the drought conditions during the month corresponding to the minimum value in Fig. 3.19. Although both indices indicate severe drought across much of the region, key differences emerge. SPI-12 depicts a less heterogeneous drought pattern, with particularly severe conditions in northern Italy and parts of southern France (where values drop below -2.5), but severe drought conditions appear, for example, in Spain and Sicily, suggesting a more spatially variable impact. In contrast, SPEI-12 presents a more uniformly distributed and intense drought, with a larger portion of the region experiencing extreme dryness. Hence, while precipitation deficits were clearly significant, the additional drying effect from high temperatures and increased evapotranspiration likely exacerbated the drought's severity, leading to a stronger and wider signal in SPEI-12.

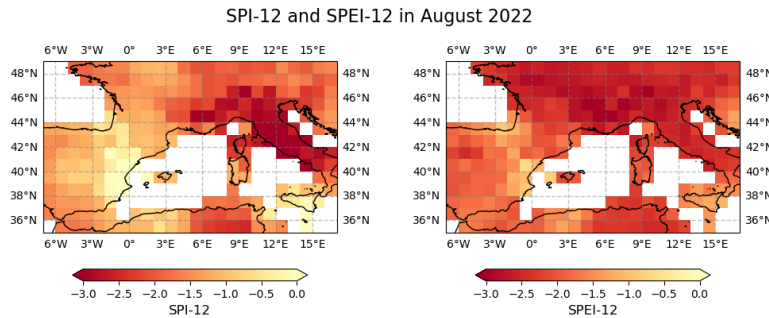


Figure 3.20: Spatial distribution of SPI-12 (left) and SPEI-12 (right) in August 2005, the month corresponding to the minimum value identified in Fig. 3.19.

Fig. 3.21 presents the characterization of precipitation anomalies over the study area in Fig. 3.19 in terms of SOM nodes. The distributions are broadly similar to the Iberian case (Fig. 3.15), which is not surprising given the similar spatial domain. However, unlike the Iberian drought's case, node (4,6) does not exhibit the strongest mean anomaly. Fig. 3.21b reveals a different pattern for positive anomalies: node (0,6) is no longer associated with the

most intense positive mean anomaly Fig. 3.21.c. Instead, other nodes, such as (0,0) and (4,0), exhibit higher positive precipitation anomalies despite occurring less frequently.

Comparing Fig. 3.21.a and Fig. 3.21.b the spatial distributions of negative and positive anomalies appear asymmetric. Negative anomalies are more evenly distributed across the array, whereas positive anomalies are more localized, with concentrations in the upper right and upper left corners. Finally, panel 3.21.c aligns closely with both Fig. 3.21.a and Fig. 3.21.b, although the relationship is less direct between Fig. 3.21.a and Fig. 3.21.c. This reinforces the idea that positive precipitation anomalies tend to be more spatially structured, whereas negative anomalies are more broadly distributed across different circulation patterns.

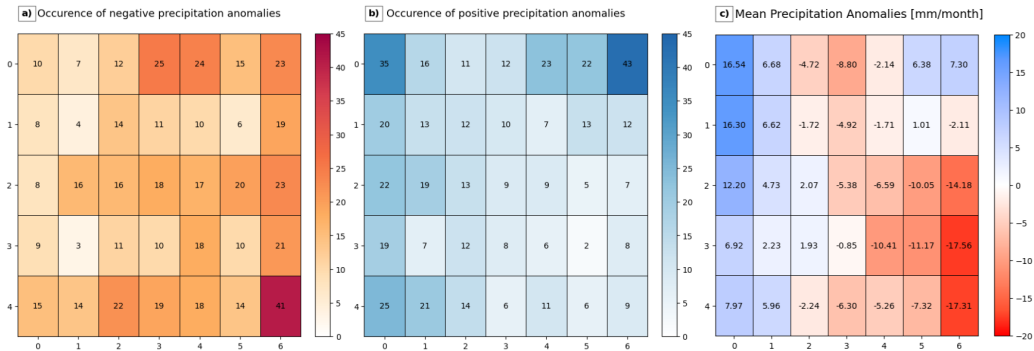


Figure 3.21: Occurrences of a) negative precipitation anomalies and b) positive precipitation anomalies over southern Europe study area, and c) mean precipitation anomaly associated with each SOM node.

The characterization of dry, wet, and neutral nodes in terms of a) SPI-1 and b) SPEI-1 (Fig. 3.22) also exhibit a similar pattern compared to the previous case and consistency with Fig. 3.21.c. SPI-1 identifies fewer dry nodes and more wet nodes, whereas SPEI-1 detects a significantly higher number of dry nodes and fewer wet ones. This discrepancy highlights the fact that SPEI-1 captures more monthly circulation patterns associated with dry conditions,

increasing the number of classified dry nodes while reducing wet classifications.

Given that this region, similarly to the Iberian Peninsula, experiences a Mediterranean climate characterized by hot, dry summers, the similarities between the two cases are not surprising. The persistent high temperatures and strong evapotranspiration rates in the region likely explain why SPEI-1 systematically detects a more widespread and severe drought compared to SPI-1, reinforcing the importance of considering atmospheric demand in drought assessments.




Figure 3.22: Characterization of wet (blue), dry (red), and neutral (yellow) SOM nodes based on a) SPI-1 and b) SPEI-1 for southern Europe study area.

Figure 3.23 shows the activation frequencies of dry, wet, and neutral nodes, represented in terms of a) SPI-1 and b) SPEI-1, compared to the reference variability. In this case, for SPI-1, the activation frequency of the dry (wet) nodes falls out of the IQR, indicating a notable increase (reduction) during the event. The activation of neutral nodes is significantly above the median, even exceeding the whiskers, implying that neutral conditions were far more persistent than usual. This suggests that under SPI-1 the drought was driven more by a strong persistence of dry and neutral configurations at the expense of wet patterns.

The characterization based on SPEI-1 is, however, markedly different. The

activation frequency for dry nodes is an extreme outlier, well above the typical distribution, highlighting the exceptional nature of the drought when considering both precipitation deficits and increased evaporative demand. Meanwhile, the wet node frequency is near zero, aligning with the lower tail of the reference distribution, further emphasizing the rarity of wet conditions during this period. This comparison suggests that, unlike previous drought events analyzed, SPEI-1 more clearly captures the exceptionality of the 2022-2023 drought, likely due to the compounded effect of reduced precipitation and abnormally high temperatures.




Figure 3.23: Frequency distribution for dry, wet, and neutral nodes considering all the periods with the same start (end) month. Red dots represent the observed frequencies for the event. The box spans from the 25th to the 75th percentile, with a line at the median. Whiskers extend from $25Q - 1.5 \times IQR$ to $75Q + 1.5 \times IQR$, where $IQR = 75Q - 25Q$.

Table 3.3 presents the observed activation frequencies, their corresponding percentiles, and the associated probabilities. The results align with the conclusions drawn in the previous paragraph. Under SPI-1, the activation fre-

quency of dry nodes appears relatively typical, with an associated probability of 14.29%, indicating that dry conditions during the event were rare but not exceptional compared to historical variability. Similar conclusions can be drawn from the probability for wet nodes (9.73%). However, the neutral nodes exhibit a notably high percentile (95.24) and a low probability (4.76%), indicating that their activation was unusually frequent compared to historical records. This reinforces the idea that, under SPI-1, the drought was characterized more by a reduction in wet conditions and an increased presence of neutral configurations rather than an extreme dominance of dry nodes.

In contrast, under SPEI-1, the activation frequency of dry nodes is extreme, reaching a percentile of 100, highlighting the exceptional and unprecedented nature of the dry conditions when accounting for both precipitation deficits and increased evaporative demand. This is further supported by the extremely low probability (0.00%), confirming that such a high frequency of dry conditions had never occurred in the reference dataset. Conversely, wet and neutral nodes exhibit very low percentiles (0.00 and 1.43, respectively), highlighting their complete absence during the event.

	SPI-1			SPEI-1		
	Freq. (%)	Pctl	Prob (%)	Freq. (%)	Pctl	Prob (%)
Dry	37.04	85.71	14.29	85.19	100.00	0.00
Wet	25.92	9.73	9.73	3.70	0.00	0.00
Neutral	37.04	95.24	4.76	11.11	1.43	1.43

Table 3.3: Observed frequencies, percentiles, and probabilities for SPI-1 and SPEI-1 during the 2022-2023 drought.

Chapter 4

Discussion and Conclusions

In this chapter, we will discuss the results presented in the previous chapter in the context of the recent scientific literature. We will examine both the benefits and potential drawbacks of the approach, highlighting its strengths as well as any limitations that may affect the interpretation of the findings. Additionally, we will consider possible improvements and alternative methods that could refine the analysis. Finally, we will explore future perspectives, discussing how the framework we defined can be expanded for further studies and potentially applied in the field.

4.1 Summary and discussion of the results

By applying a 7×5 Self-Organizing Map, we identified the predominant atmospheric patterns over the Euro-Atlantic domain and used them to characterize three major drought events using two drought indicators, SPI-1 and SPEI-1.

For the 1974-1975 drought (Parry et al. (2012)), SPI-1 identified 11 wet patterns, 9 dry patterns, and the rest as neutral out of 35 total patterns, while SPEI-1 classified 6 as wet, 10 as dry, and the rest as neutral. Overall, the drought seems to be primarily driven by an above-average activation of

dry nodes, while the activation of wet nodes is substantially below-average (10.53% for SPI and 5.26% for SPEI). The rarity of wet nodes is particularly striking, with probabilities of only 1.64% (SPI-1) and 0.35% (SPEI-1), highlighting an exceptional scarcity of precipitation-bearing patterns. The main difference between SPI-1 and SPEI-1 emerges in the frequency of activation for wet nodes, which is half in SPEI-1 (5.26%) compared to SPI-1 (10.53%), likely reflecting the additional influence of evapotranspiration.

For the 2004-2005 Iberian drought (García-Herrera et al. (2007)), the differences between SPI and SPEI are more pronounced, especially in the identification of dry and wet nodes. While SPI-1 classifies 8 dry nodes and 7 wet nodes, SPEI-1 identifies 16 dry nodes and only 5 wet nodes, highlighting the region's sensitivity to evapotranspiration effects (Vicente-Serrano and López-Moreno (2008); González-Hidalgo et al. (2018)). These differences are also reflected in the frequency and probability values. Under SPI-1, the drought appears to be driven by an unusually high activation of neutral nodes (50.00%) at the expense of wet nodes (16.67%), with probability values indicating that such persistence of neutral patterns was virtually unprecedented in the dataset (0.14%). In contrast, SPEI-1 attributes the drought primarily to the high activation of dry nodes (41.67%), which are associated with a probability of 6.25%, while the occurrence of neutral nodes (21.43%) is far less exceptional.

Lastly, for the 2022-2023 drought (Bevacqua et al. (2024); Faranda et al. (2023)), the results strongly reflect the unprecedented severity of the event. The classification of nodes remains almost identical to the previous case, given the similar geographical domain, with SPI-1 identifying 8 dry and 7 wet nodes, while SPEI-1 classifies 17 dry and 5 wet nodes. However, the most remarkable aspect is the activation frequency under SPEI-1: dry nodes were active for the vast majority of the time (85.19%), at the expense of both wet (3.70%) and neutral nodes (11.11%). In practical terms, wet nodes were almost entirely absent during the event. The probability values further

highlight the extremeness of the drought by showing that dry nodes had an associated 0% probability, indicating that such high activation has never occurred in the observational dataset ranging over the period 1940-2023. The probability for wet nodes activation was also exceptionally low (0%), reinforcing the unprecedented nature of the event.

These results demonstrate that the methodology we developed is not only effective in identifying extreme events, but also in providing new insights into their underlying dynamics. First, it offers a continuous representation of the dominant atmospheric patterns over the Euro-Atlantic domain, improving upon traditional approaches. Second, it allows for a quantification evaluation of the probability of occurrence of dry, wet, and neutral circulation patterns, offering a deeper insight into the dynamical drivers of the event.

As far as the methodology is concerned, some critical aspects must be acknowledged. A key challenge is the selection of the number of nodes in the SOM array, as different configurations can lead to significantly different results. There is no universal rule for determining the optimal number of nodes; instead, it largely depends on the study's objectives and the desired level of detail. This remains an open area for further experimentation, as different SOM configurations could reveal additional insights into the complex dynamics behind droughts. However, we expect that the key point of this approach, i.e., the identification of dry, wet, and neutral nodes and their aggregated probability of occurrence during an extreme drought event, is pretty insensitive to the change of SOM dimension.

Another potential limitation lies in the choice of SPI-1 and SPEI-1 as criteria for defining dry, wet, and neutral nodes. While these indices are widely used in drought studies, alternative approaches exist. For instance, [Odoulami et al. \(2020\)](#) and [Harrington et al. \(2016\)](#) classified atmospheric patterns by analyzing the broader circulation dynamics over their domain of interest. Instead of relying solely on drought indices, they assigned labels to patterns

based on their role in regional climate variability. Specifically, they considered the positioning of key synoptic features, such as westerlies, subtropical anticyclones, and low-pressure systems, and identified patterns that actively contributed to blocking as "dry" and those related to rain-bearing systems as "wet". Future research could explore hybrid methods that integrate multiple criteria to enhance classification robustness, such as incorporating key synoptic features that influence the spatial domain. A third point requiring some attention is the time available for analysis, which in our case was 83 years (1940-2023). Having reanalysis and longer precipitation datasets, such as 20CRU3, would provide a more robust historical perspective. Despite these limitations, the methodology remains a powerful tool for studying droughts from a new perspective and lays the foundation for attribution studies.

The next step for this study would be to conduct an extreme event attribution analysis (Otto (2017)) to assess the influence of anthropogenic climate change on these drought events. Our method not only provides insights into how human-induced warming affects drought likelihood but, more importantly, helps evaluate its impact on atmospheric dynamical drivers. Quantifying this influence is particularly challenging, and few studies have successfully isolated clear anthropogenic signals in circulation patterns. While thermodynamic effects, such as those driven by warming, are relatively easier to detect, the impact of human-induced changes on purely dynamical atmospheric variables remains highly uncertain (Bevacqua et al. (2024)).

Our approach would be especially valuable for analyzing multi-year droughts (Pascale and Ragone (2025)), which involves the synergy of different atmospheric drivers to maintain dry conditions over a long time period. While standard model comparisons remain a complementary method, applying such techniques to drought studies often relies on the assumption that a specific drought measure, such as soil moisture, can be consistently simulated by climate models over time and is fully representative of the underlying mech-

anisms driving event severity. In contrast, the SOM-based approach focuses solely on changes in atmospheric circulation patterns throughout the event, providing a unique perspective on drought attribution that is less dependent on a single drought metric.

The proposed framework, inspired by [Odoulami et al. \(2023\)](#), would aim to assess the impact of anthropogenic climate change by comparing climate model simulations under different forcing scenarios. It consists of the following steps:

- 1) Estimate the frequency of activation for dry, wet, and neutral nodes during the event.
- 2) Obtain the same estimate in a counterfactual climate, where human-induced climate change is absent, and in a factual climate, where it is present.
- 3) Assess whether anthropogenic climate change has significantly altered the probability of having such an extreme event.

To achieve this, we employed simulations from the atmospheric general circulation model ECHAM5.4 ([Hoell et al. \(2017\)](#)), obtained from the Atmosphere Model Intercomparison Project (AMIP; [Hoell et al. \(2017\)](#)), freely available at <https://ps1.noaa.gov/repository/entry/show>. These simulations include a large ensemble of atmosphere-only experiments under two forcing scenarios: all or actual forcings (ALL) representing the factual world, and natural-only forcing (NAT) as the counterfactual. In the NAT simulations, (i) observed sea surface temperatures are detrended and adjusted to 1880 mean conditions, (ii) sea ice follows a repeating seasonal cycle representative of 1979–1989 conditions, and (iii) greenhouse gas concentrations and ozone levels are set to their 1880 values.

Applying the SOM model to these datasets would allow us to redefine and analyze the dry, wet, and neutral nodes under both scenarios. By computing activation frequencies, we could quantify how anthropogenic influences

have altered the persistence and occurrence of specific atmospheric patterns, ultimately assessing whether human-induced climate change has increased or decreased the likelihood of extreme drought events. Fig. 4.1 shows the SOM-based classification of Z_{500} monthly anomalies in the NAT ECHAM5.4 ensemble. Apart from differences in the labeling and orientation of the SOM array, comparison to Fig. 3.1 shows the ability of the model to capture the main dry, wet, and neutral nodes. The next step, which goes beyond the scope of this thesis and is left for future work, would be to quantify the probabilities of occurrence of dry, wet, and neutral nodes for a specific event (e.g. for the 2022-2023 Euro-Mediterranean drought) in the NAT and ALL ensembles, and, eventually, extend these to additional large ensemble models (like, e.g. in Pascale et al. (2020)).

4.2 Conclusions

In this thesis, we examined the dynamical drivers of extreme, major drought events in Europe through a SOM-based framework. We implemented a 7×5 Self-Organizing Map to identify the most recurrent Z_{500} anomaly patterns across Europe and North Atlantic, and to correlate these patterns with the significant teleconnections impacting Europe. We found that this approach offers a continuous depiction of atmospheric states, revealing intermediate states that do not represent a single mode, but rather the superimposition of many. Most importantly, we showed how these patterns identified through the SOM affect precipitation, near-surface temperature, and potential evapotranspiration, which are the main variables defining droughts.

The key findings of this work are about the characterization of three major extreme European drought events (1975-1976 UK drought, 2004-2005 Iberian drought, and 2022-2023 Euro-Mediterranean drought). While few other studies have employed similar methodologies (Odoulami et al., 2020, 2023), and Harrington et al. (2016)), none have specifically addressed these three events,

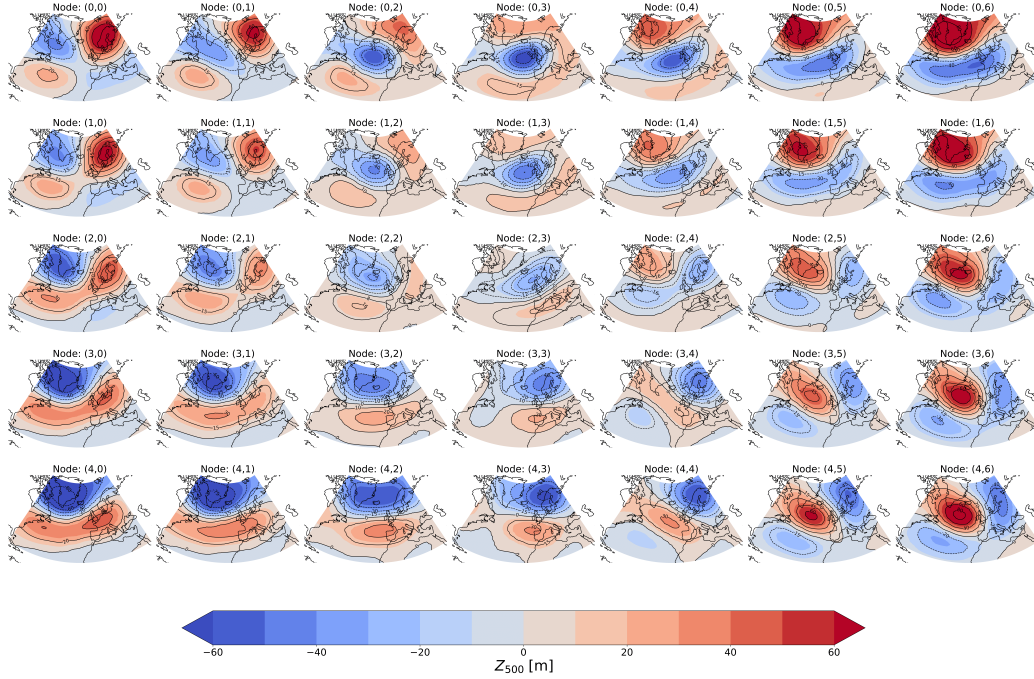


Figure 4.1: Z_{500} monthly anomalies patterns for the NAT experiment from the general circulation model ECHAM5.4 from AMIP (Hoell et al., 2017) classified using a SOM with 35 nodes.

making this analysis a novel contribution to the field. On the other hand, while these events have been extensively studied (Parry et al. (2012); García-Herrera et al. (2007); Faranda et al. (2023)), none of these focused on the dynamical drivers using the SOM approach. Our results demonstrate the effectiveness of the SOM approach in recognizing the exceptional nature of these events. Specifically, we showed that during the events the patterns associated with dry conditions manifested significantly more frequently than rain-bearing ones, providing a new insight into the underlying dynamic driving the event. Finally, the combined use of SPI and SPEI clearly differentiates the dynamical and thermodynamical drivers of the droughts. For example,

in the case of the recent drought in 2022-2023, we proved that high temperatures contributed significantly to the severity of the event, leading to unprecedented drying conditions. In future work, the application of such an approach to a large ensemble of climate models will allow for an attribution of the driving mechanism of major European droughts.

Appendix A

Appendix

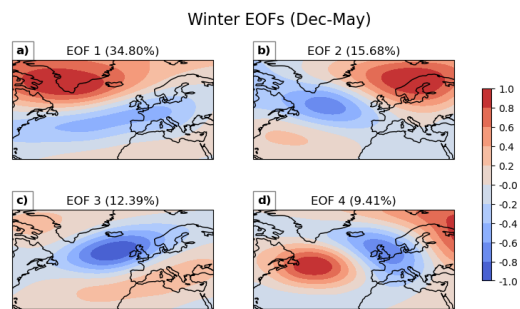


Figure A.1: Z_{500} EOFs during extended winter months.

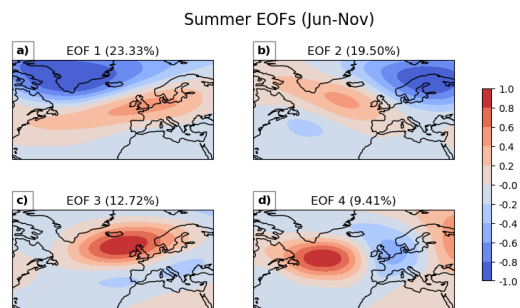


Figure A.2: Z_{500} EOFs during extended summer months.

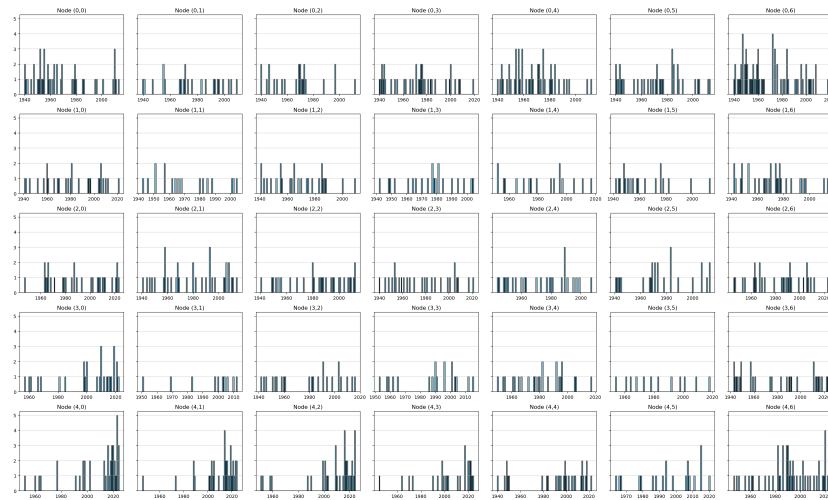


Figure A.4: Yearly frequency for every node



Figure A.3: Monthly frequency for every node




Figure A.5: Same analysis computed only in a) extended winter (Dec-May) and b) extended summer (Jun-Nov).




Figure A.6: Same analysis computed only in a) extended winter (Dec-May) and b) extended summer (Jun-Nov).

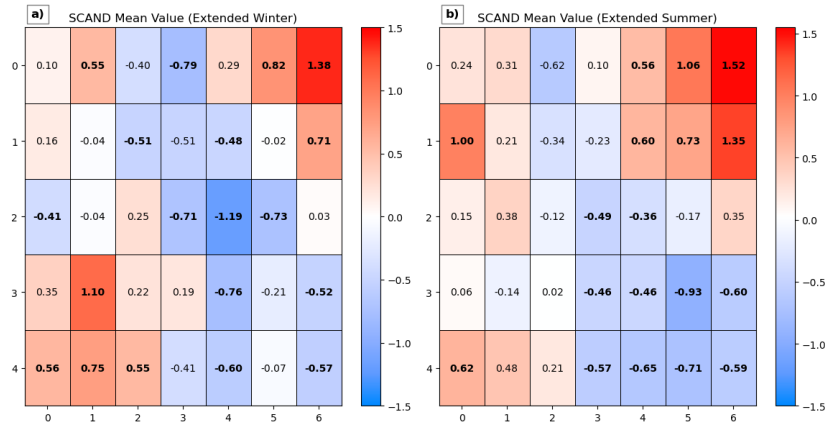


Figure A.7: Same analysis computed only in a) extended winter (Dec-May) and in b) extended summer (Jun-Nov).

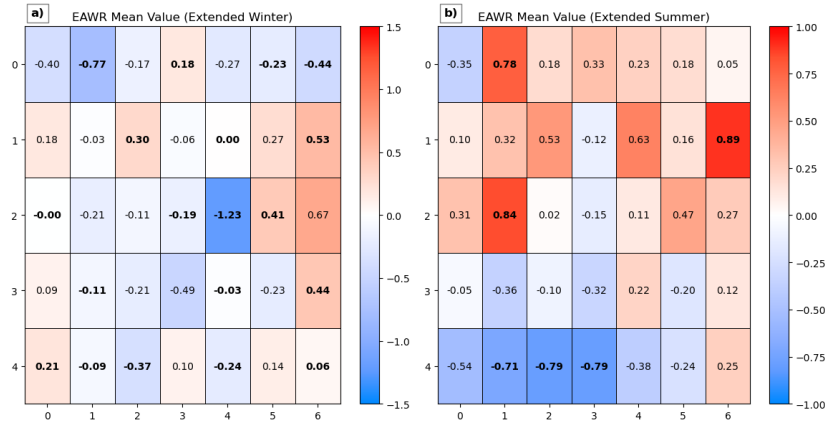


Figure A.8: Same analysis computed only in a) extended winter (Dec-May) and b) extended summer (Jun-Nov).

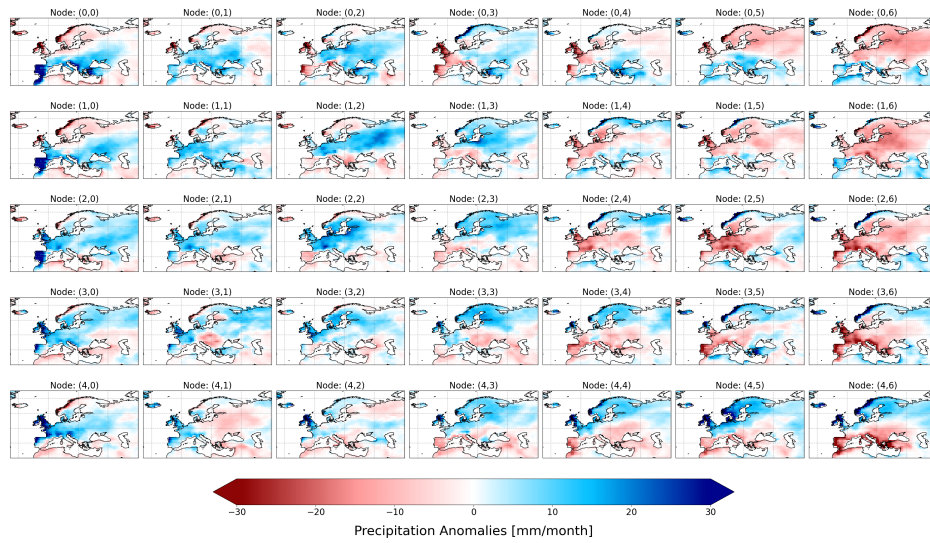


Figure A.9: Composites of precipitation anomalies for each SOM node.

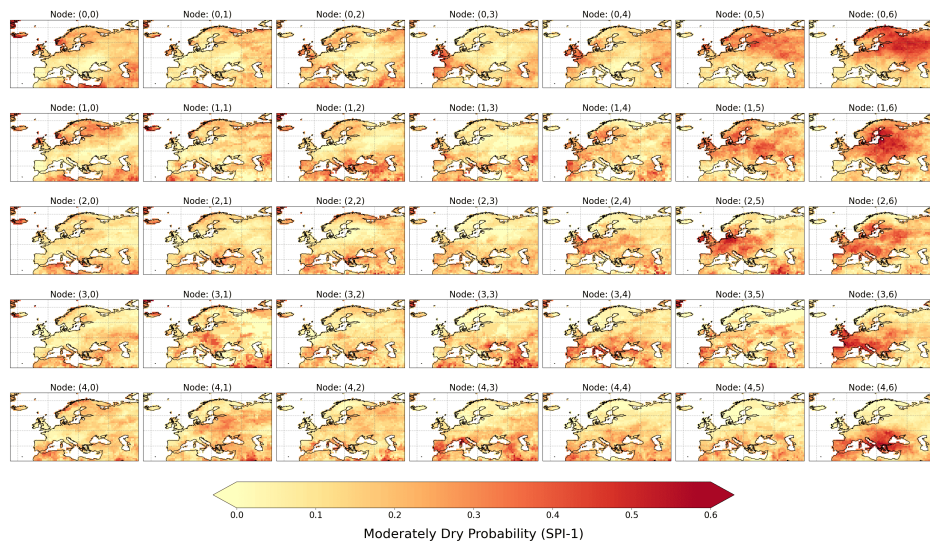


Figure A.10: Maps of moderately dry probability for SPI-1.

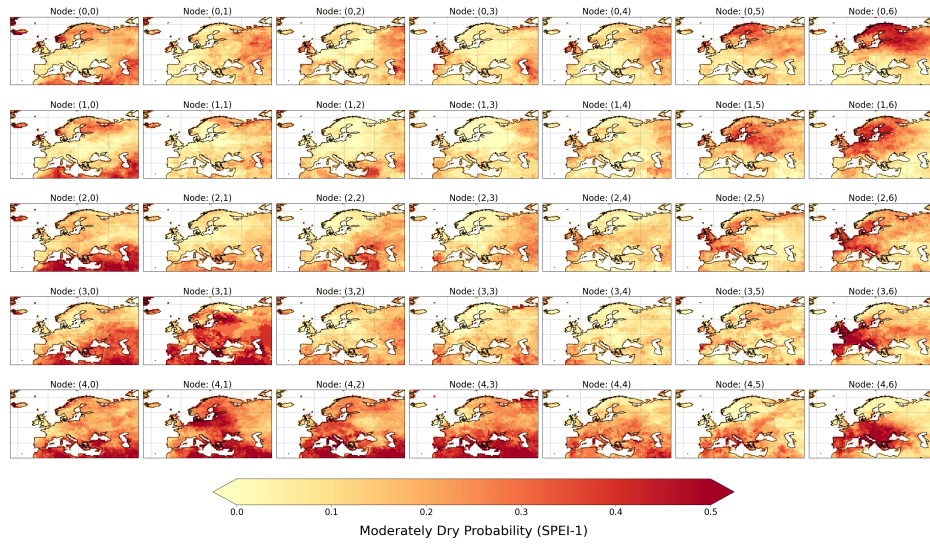


Figure A.11: Maps of moderately dry probability for SPEI-1.

Bibliography

- Barnston, A. G. and Livezey, R. E. (1987). Classification, Seasonality and Persistence of Low-Frequency Atmospheric Circulation Patterns. *Monthly Weather Review*, 115(6):1083–1126. Publisher: American Meteorological Society Section: Monthly Weather Review.
- Baronetti, A., Dubreuil, V., Provenzale, A., and Fratianni, S. (2022). Future droughts in northern Italy: high-resolution projections using EURO-CORDEX and MED-CORDEX ensembles. *Climatic Change*, 172(3):22.
- Baronetti, A., Gonzalez-Hidalgo, J., Vicente-Serrano, S., Acquaotta, F., and Simona, F. (2020). A weekly spatio-temporal distribution of drought events over the Po Plain (North Italy) in the last five decades. *International Journal of Climatology*, 40.
- Bellman, R. (1961). *Adaptive Control Processes: A Guided Tour*. Princeton University Press.
- Bevacqua, E., Rakovec, O., Schumacher, D., Kumar, R., Thober, S., Samaniego, L., Seneviratne, S., and Zscheischler, J. (2024). Direct and lagged climate change effects intensified the 2022 european drought. *Nature Geoscience*, 17:1100–1107.
- Bordi, I. and Sutera, A. (2002). An analysis of drought in Italy in the last fifty years. *Il Nuovo Cimento C*, 25.

- Cavazos, T. (2000). Using Self-Organizing Maps to Investigate Extreme Climate Events: An Application to Wintertime Precipitation in the Balkans. *Journal of Climate - J CLIMATE*, 13:1718–1732.
- Dai, A. (2011a). Characteristics and trends in various forms of the Palmer Drought Severity Index during 1900–2008. *Journal of Geophysical Research: Atmospheres*, 116(D12). _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2010JD015541>.
- Dai, A. (2011b). Drought Under Global Warming: A Review. *Wiley Interdisciplinary Reviews: Climate Change*, 2:45–65.
- Esteban-Parra, M. J., Rodrigo, F. S., and Castro-Diez, Y. (1998). Spatial and temporal patterns of precipitation in Spain for the period 1880–1992. *International Journal of Climatology*, 18(14):1557–1574. _eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/%28SICI%291097-0088%2819981130%2918%3A14%3C1557%3A%3AAID-JOC328%3E3.0.CO%3B2-J>.
- Faranda, D., Pascale, S., and Bulut, B. (2023). Persistent anticyclonic conditions and climate change exacerbated the exceptional 2022 European-Mediterranean drought. *Environmental Research Letters*.
- García-Herrera, R., Hernández, E., Barriopedro, D., Paredes, D., Trigo, R. M., Trigo, I. F., and Mendes, M. A. (2007). The Outstanding 2004/05 Drought in the Iberian Peninsula: Associated Atmospheric Circulation. *Journal of Hydrometeorology*, 8(3):483–498. Publisher: American Meteorological Society Section: Journal of Hydrometeorology.
- Garrido-Perez, J., Vicente-Serrano, S., Barriopedro, D., García-Herrera, R., Trigo, I., and Beguería, S. (2024). Examining the outstanding Euro-Mediterranean drought of 2021–2022 and its historical context. *Journal of Hydrology*, 630:130653.

- Giorgi, F. (2006). Climate change hot-spots. *Geophysical Research Letters*, 33(8). [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1029/2006GL025734](https://onlinelibrary.wiley.com/doi/pdf/10.1029/2006GL025734).
- González-Hidalgo, J. C., Vicente-Serrano, S. M., Peña-Angulo, D., Salinas, C., Tomas-Burguera, M., and Beguería, S. (2018). High-resolution spatio-temporal analyses of drought episodes in the western Mediterranean basin (Spanish mainland, Iberian Peninsula). *Acta Geophysica*, 66(3):381–392.
- Hargreaves, G. and Samani, Z. (1985). Reference crop evapotranspiration from temperature. *Applied Engineering in Agriculture*, 1:96–99.
- Harrington, L. J., Gibson, P. B., Dean, S. M., Mitchell, D., Rosier, S. M., and Frame, D. J. (2016). Investigating event-specific drought attribution using self-organizing maps. *Journal of Geophysical Research: Atmospheres*, 121(21):12,766–12,780. [_eprint: https://onlinelibrary.wiley.com/doi/pdf/10.1002/2016JD025602](https://onlinelibrary.wiley.com/doi/pdf/10.1002/2016JD025602).
- Harris, I., Osborn, T. J., Jones, P., and Lister, D. (2020). Version 4 of the CRU TS monthly high-resolution gridded multivariate climate dataset. *Scientific Data*, 7(1):109. Publisher: Nature Publishing Group.
- Hersbach, H., Bell, B., Berrisford, P., Hirahara, S., Horányi, A., Muñoz Sabater, J., Nicolas, J., Peubey, C., Radu, R., Schepers, D., Simmons, A., Soci, C., Abdalla, S., Abellan, X., Balsamo, G., Bechtold, P., Biavati, G., Bidlot, J., Bonavita, M., and Thépaut, J.-N. (2020). The ERA5 global reanalysis. *Quarterly Journal of the Royal Meteorological Society*, 146.
- Hewitson, B. and Crane, R. (2002). Self-Organizing Maps: Applications to synoptic climatology. *Climate Research*, 22:13–26.
- Hoell, A., Hoerling, M., Eischeid, J., Quan, X. W., and Liebmann, B. (2017).

- Reconciling theories for human and natural attribution of recent east africa drying. *Journal of Climate*, 30:1939–1957.
- IPCC (2023). *Climate Change 2023: Synthesis Report. Contribution of Working Groups I, II and III to the Sixth Assessment Report of the Intergovernmental Panel on Climate Change*. IPCC, Geneva, Switzerland.
- Johnson, N. (2013). How many ENSO flavors can we distinguish? *Journal of Climate*, 26:4816–4827.
- Johnson, N. C., Feldstein, S. B., and Tremblay, B. (2008). The Continuum of Northern Hemisphere Teleconnection Patterns and a Description of the NAO Shift with the Use of Self-Organizing Maps. *Journal of Climate*, 21(23):6354–6371. Publisher: American Meteorological Society Section: Journal of Climate.
- Kalnay, E., Kanamitsu, M., Kistler, R., Collins, W., Deaven, D., Gandin, L., Iredell, M., Saha, S., White, G., Woollen, J., Zhu, Y., Chelliah, M., Ebisuzaki, W., Higgins, W., Janowiak, J., C, K., Ropelewski, C., Wang, J., and Leetmaa, A. (1996). The NMC/NCAR 40-year reanalysis project. *Bull Am Meteorol Soc. Bulletin of the American Meteorological Society*, 77.
- Kingston, D. G., Stagge, J. H., Tallaksen, L. M., and Hannah, D. M. (2015). European-Scale Drought: Understanding Connections between Atmospheric Circulation and Meteorological Drought Indices. *Journal of Climate*, 28(2):505–516. Publisher: American Meteorological Society Section: Journal of Climate.
- Li, X., He, B., Quan, X., Liao, Z., and Bai, X. (2015). Use of the Standardized Precipitation Evapotranspiration Index (SPEI) to Characterize the Drying Trend in Southwest China from 1982–2012. *Remote Sensing*, 7:10917–10937.

- Marsh, T., Cole, G., and Wilby, R. (2007). Major droughts in England and Wales, 1800–2006. *Weather*, 62(4):87–93. eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1002/wea.67>.
- McKee, T. B., Doesken, N. J., and Kleist, J. (1993). The relationship of drought frequency and duration to time scales. *Eighth Conference on Applied Climatology, 17-22 January 1993, Anaheim, California*.
- Mishra, A. K. and Singh, V. P. (2010). A review of drought concepts. *Journal of Hydrology*, 391(1):202–216.
- Monteith, J. L. (1965). Evaporation and environment. *Symposia of the Society for Experimental Biology*, 19:205–234. Publisher: Cambridge University Press (CUP) Cambridge.
- Odoulami, R., Wolski, P., and New, M. (2020). A SOM-based analysis of the drivers of the 2015–2017 Western Cape drought in South Africa. *International Journal of Climatology*, 41.
- Odoulami, R., Wolski, P., and New, M. (2023). Attributing the driving mechanisms of the 2015–2017 drought in the Western Cape (South Africa) using self-organising maps. *Environmental Research Letters*, 18.
- Otto, F. E. (2017). Attribution of weather and climate events. *Annual Review of Environment and Resources*, 42:627–646. First published as a Review in Advance on July 26, 2017.
- Palmer, W. C. (1965). *Meteorological Drought*. U.S. Department of Commerce, Weather Bureau.
- Parry, S., Hannaford, J., Lloyd-Hughes, B., and Prudhomme, C. (2012). Multi-year droughts in Europe: Analysis of development and causes. *Hydrology Research*, 43:689–706.

- Pascale, S., Kapnick, S. B., Delworth, T. L., and Cooke, W. F. (2020). Increasing risk of another Cape Town “Day Zero” drought in the 21st century. *Proceedings of the National Academy of Sciences*, 117(47):29495–29503. Publisher: Proceedings of the National Academy of Sciences.
- Pascale, S. and Ragone, F. (2025). Widespread multi-year droughts in Italy: identification and causes of development. *International Journal of Climatology*. In press.
- Reusch, D., Alley, R., and Hewitson, B. (2005). Relative Performance of Self-Organizing Maps and Principal Component Analysis in Pattern Extraction from Synthetic Climatological Data. *Polar Geography*, 29:188–212.
- Reusch, D. B., Alley, R. B., and Hewitson, B. C. (2007). North Atlantic climate variability from a self-organizing map perspective. *Journal of Geophysical Research: Atmospheres*, 112(D2). eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2006JD007460>.
- Ritter, H. and Kohonen, T. (1989). Self-organizing semantic maps. *Biological Cybernetics*, 61(4):241–254.
- Rodda, J. and Marsh, T. (2011). The 1975-76 Drought: A contemporary and retrospective review. Technical report, Centre for Ecology & Hydrology (CEH).
- Rousi, , Anagnostopoulou, C., Tolika, K., and Maheras, P. (2015). Representing teleconnection patterns over Europe: A comparison of SOM and PCA methods. *Atmospheric Research*, 152:123–137.
- Schulzweida, U. (2023). Cdo user guide.
- Skific, N. and Francis, J. (2012). Self-Organizing Maps: A Powerful Tool for the Atmospheric Sciences. In Johnsson, M., editor, *Applications of Self-Organizing Maps*. InTech.

- Sousa, P. M., Trigo, R. M., Aizpurua, P., Nieto, R., Gimeno, L., and Garcia-Herrera, R. (2011). Trends and extremes of drought indices throughout the 20th century in the Mediterranean. *Natural Hazards and Earth System Sciences*, 11(1):33–51. Publisher: Copernicus GmbH.
- Spinoni, J., Naumann, G., Vogt, J., and Barbosa, P. (2015a). European drought climatologies and trends based on a multi-indicator approach. *Global and Planetary Change*, 127:50–57.
- Spinoni, J., Naumann, G., Vogt, J. V., and Barbosa, P. (2015b). The biggest drought events in Europe from 1950 to 2012. *Journal of Hydrology: Regional Studies*, 3:509–524.
- Thorntwaite, M. (1951). The Role of Evapotranspiration in Climate. pages 16–39.
- Van der Wiel, K., Batelaan, T., and Wanders, N. (2022). Large increases of multi-year droughts in north-western Europe in a warmer climate. *Climate Dynamics*, 60:1–20.
- Vicente-Serrano, S. M., Beguería, S., and 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.
- Vicente-Serrano, S. M. and López-Moreno, J. I. (2008). Nonstationary influence of the North Atlantic Oscillation on European precipitation. *Journal of Geophysical Research: Atmospheres*, 113(D20). eprint: <https://onlinelibrary.wiley.com/doi/pdf/10.1029/2008JD010382>.
- Wells, N., Goddard, S., and Hayes, M. (2004). A Self-Calibrating Palmer Drought Severity Index. *Journal of Climate - J CLIMATE*, 17:2335–2351.
- Wilhite, D. (2000). Drought as a natural hazard: Concepts and definitions. *Drought, a Global Assessment*, 1.

Willhite, D. and Glantz, M. (1985). Understanding: the Drought Phenomenon: The Role of Definitions. *Water International - WATER INT*, 10:111-120.