Targeted weather regimes identify circulation patterns behind Western European summer heat extremes and trends

Julianna Carvalho-Oliveira[1,*], Fiona R. Spuler [1,2] and Marlene Kretschmer[1]

[1]Leipzig Institute for Meteorology, Leipzig University, Leipzig, Germany
[2]Department of Meteorology, University of Reading, Reading, UK
*Author to whom any correspondence should be addressed.
Contact: julianna.carvalho_oliveira@uni-leipzig.de

Funding statement: The authors acknowledge funding from the Horizon Europe project EXPECT (Towards an Integrated Capability to Explain and Predict Regional Climate Changes), Grant Agreement No. 101137656. F.S. further acknowledges funding from the Advancing the Frontiers of Earth System Prediction (AFESP) Doctoral Training Programme at the University of Reading.

Journal Name

Crossmark

RECEIVED
dd Month yyyy

REVISED
dd Month yyyy

RESEARCH LETTER

Targeted weather regimes identify circulation patterns behind Western European summer heat extremes and trends

Julianna Carvalho-Oliveira[1,*], Fiona R. Spuler [1,2] and Marlene Kretschmer[1]

[1]Leipzig Institute for Meteorology, Leipzig University, Leipzig, Germany [2]Department of Meteorology, University of Reading, Reading, UK [*]Author to whom any correspondence should be addressed.

E-mail: julianna.carvalho_oliveira@uni-leipzig.de

Keywords: heat extremes, weather regimes, temperature trends, machine learning, variational autoencoders

1 Abstract

Western European heat extremes have intensified in recent decades, with their rate of warming outpacing the global mean. Against this general human-induced warming trend, understanding the circulation patterns that drive such heat extremes is crucial. Weather-regime (WR) approaches have been widely used to characterise large-scale circulation variability; however, conventional classifications are not optimised to identify the dynamical drivers of extremes. Here we apply a novel targeted machine learning-based approach, the regression mixture-model variational autoencoder (RMM-VAE), to characterise summer weather regimes relevant for Western European temperatures. We compare its performance against the standard, non-targeted k-means approach and find RMM-VAE to yield dynamically coherent regimes while being more informative of (extreme) temperatures in the target region.

Our analysis identifies a southerly-flow regime that accounts for the vast majority of heatwave days, while k-means disperses them across multiple regimes. Moreover, the seasonal frequency of this impact-relevant regime, combined with global mean temperature, explains a large fraction of interannual variability in both mean ($R^2 = 0.84$) and extreme summer temperatures (90th percentile; $R^2 = 0.65$), with predictive skill persisting out-of-sample tests. Finally, this simple regression model allows us to attribute 34% of total summer warming in Western Europe and about 70% of the observed "excess" warming (relative to the global mean) to an observed increase in the identified southerly flow circulation patter, which we quantify to be largely forced.

Our results demonstrate that targeted weather regime approaches can sharpen the link between circulation and surface extremes, offering attribution of regional warming. Furthermore, the identified regimes provide interpretable predictors with potential for improving seasonal forecasts and climate risk assessments.

2 Introduction

Heat extremes in Western Europe have increased in frequency and intensity over the last decades, with associated impacts on health [Chen et al., 2025], mortality [Campbell et al., 2018], agriculture [Lesk et al., 2016], energy systems [Brás et al., 2023], and ecosystems [Von Buttlar et al., 2018]. Recent summers, such as in 2018 and 2022, have highlighted both their devastating consequences and the challenges of anticipating their occurrence [Rousi et al., 2023, Li et al., 2024]. While rising global mean temperature (GMT) due to anthropogenic climate change sets the thermodynamic background, recent studies indicate that changes in large-scale circulation play a crucial role in observed trends in European temperatures, contributing to the region's "excess warming" relative to the global mean [Rousi et al., 2022, Singh et al., 2023, Dong and Sutton, 2025, Pfleiderer et al., 2025]. More precisely, observations indicate more frequent and persistent anticyclonic conditions over western and central Europe [Horton et al., 2015], a weakened summertime storm track and zonal flow [Coumou et al., 2015, Lehmann and Coumou, 2015], and episodes of enhanced meridional jet excursions and quasi-stationary wave activity [Coumou et al., 2014, Kornhuber et al., 2017], all of which favour warm-air advection and subsidence over the continent. State-of-the-art climate models, however, struggle to reproduce these observed circulation trends, suggesting that the dynamical contribution to observed and future warming might be underestimated [Van Oldenborgh et al., 2022, Vautard et al., 2023].

An open question is how much these observed circulation changes reflect internal variability versus a forced response of the climate system. Internal variability is substantial across

sub-monthly to multi-decadal scales, giving the observational record a low signal-to-noise ratio [Deser et al., 2012]. Different modes of variability often overlap and project onto similar circulation patterns, so regional trends may partly reflect the timing and phasing of internal variability rather than an externally forced shift, which complicates attribution of circulation trends [Deser and Phillips, 2023]. A large share of this variability stems from random stochastic processes in the atmospheric circulation, which operates across timescales [Deser et al., 2012], and from coupled ocean–atmosphere modes such as El Niño–Southern Oscillation (ENSO), and Atlantic multi-decadal variability (Zhang et al. [2019]) with limited predictability. Over the short instrumental record these fluctuations can obscure or resemble forced regional signals, further complicating detection and attribution.

To isolate the (forced) dynamical component of long-term temperature trends, different approaches have been applied. These include quantifying trends in simple box averages of geopotential heights (e.g. [Dong and Sutton, 2025]), and dynamical adjustment (e.g. [Vautard et al., 2023, Sippel et al., 2019, Van Oldenborgh et al., 2022]), which partition trends into thermodynamic and circulation-induced components. More recently, methods based on deep learning or creating nudged-circulation simulations fixing observed large-scale flows in pre-industrial conditions have also been proposed (for a method overview see Pfleiderer et al. [2025]). Observed trend studies often emphasise increases in southerly-flow regimes over Europe (e.g. [Vautard et al., 2023]) and more generally detect long-term shifts in circulation via European geopotential height indices [Dong and Sutton, 2025] or broader Northern Hemisphere wave-train patterns [Teng et al., 2022], contributing to long-term summer warming.

Although several recent studies have focused on attributing long-term trends in European heatwaves [Perkins-Kirkpatrick and Lewis, 2020, Felsche et al., 2024], considerable research has also examined the influence of large-scale circulation on synoptic to subseasonal timescales [Rouges et al., 2023, Happe et al., 2024]. Well documented dynamical configurations causing European heat extremes include Euro–Atlantic blocking and Atlantic lows, accompanied by a ridge over southern Europe. Such patterns are associated with a weakened zonal flow, enhanced meridional jet excursions, and favour warm-air advection toward the continent [Cassou et al., 2005, Perkins, 2015]. These circulation states are often embedded in slowly evolving, high-amplitude midlatitude wave patterns—classically described as quasi-stationary Rossby waves with hemispheric wave numbers 5–7 [Kornhuber et al., 2020] or as circumglobal wave trains with synoptic-scale wave numbers 6–8 [Di Capua et al., 2021]. The resulting mid- to upper-tropospheric anticyclones sustain poleward advection along the western flank, subsidence-driven adiabatic warming, and clearer skies, which together act to elevate surface temperatures [Domeisen et al., 2023]. At times, radiative effects and subsidence dominate over horizontal advection as the main contributors to extreme heat [Luo et al., 2022, Zschenderlein et al., 2019]. Heatwaves are further amplified by antecedent soil moisture deficits that shift surface energy partitioning from latent to sensible heat, reinforcing near-surface warming and prolonging events [Seneviratne et al., 2010, García-García et al., 2023].

A common entry point to study circulation patterns and their links to temperature is to classify circulation fields in different *weather regimes*. These refer to recurrent and persistent large-scale circulation states. In summer, the Euro–Atlantic sector has been the focus of extensive regime analyses, with studies identifying recurrent circulation states that capture blocking, ridge, North Atlantic Oscillation (NAO) and East Atlantic-like patterns [e.g. Cassou et al., 2005, Rousi et al., 2023, Carvalho-Oliveira et al., 2024]. Regimes are tied to the eddy-driven jet and help diagnose variability and link circulation to surface extremes [Madonna et al., 2017, Rouges et al., 2023, Lemburg and Fink, 2024].

The most widely applied method for computing weather regimes combines principal component analysis (PCA), a linear dimensionality reduction approach, with clustering, often using the k-means algorithm [e.g. Michelangeli et al., 1995, Grams et al., 2017, Hannachi et al., 2017]. More recently, machine learning methods that capture nonlinear, spatiotemporal features have been applied to study atmospheric circulation, including self-organizing maps and autoencoders (e.g. Carvalho-Oliveira et al. [2022], Kurihana et al. [2024]). However, while these existing approaches are designed to capture the main modes of variability over a given region, they can average out the circulation patterns that lead to particular extremes and impacts - such as heat extremes over Western Europe - making them less suitable for studying circulation drivers of local extremes. To address this, *targeted* regime approaches have been proposed. These methods generally condition regime identification on an impact variable, for example by clustering fields associated with extreme temperatures [Felsche et al., 2023, Rouges et al., 2023, Happe et al., 2024], or clustering the impact variable directly [Bloomfield et al., 2020]. Happe et al. [2024] used variational autoencoders (VAE) to classify Western European heatwaves from a large climate ensemble,

arguing that heatwaves are best viewed in probabilistic terms.

However, such targeted regimes can compromise physical consistency and predictability. To overcome this, Spuler et al. [2024] introduced the regression mixture-model VAE (RMM-VAE), which identifies targeted regimes representing the broader atmospheric phase space while being informative of the target variable. RMM-VAE combines nonlinear dimensionality reduction, clustering, and a regression task, allowing a scalar target variable alongside high-dimensional gridded circulation data. Applied to Morocco winter precipitation, it produced regimes more informative of the target variable than k-means, while maintaining predictability and physical consistency [Spuler et al., 2024, 2025].

Motivated by the findings of [Spuler et al., 2024], we here apply RMM-VAE to identify atmospheric circulation states driving variability and trends in European summer temperatures, including extreme events. Our contributions are threefold: (1) we compare RMM-VAE against k-means in identifying patterns driving Western European heat extremes; (2) we assess if the regimes also explain and predict interannual temperature variability by including seasonal regime frequency and GMT as covariate; and (3) we contextualise our results with trend attribution studies, quantifying the dynamical contribution of circulation changes to excess summer warming.

3 Data and methods

3.1 Data and preprocessing

We analyse daily ERA5 data for the period 1950–2022 [Hersbach et al., 2020], focusing on boreal summer (June–August, JJA). Unless stated otherwise, this period and seasonal selection apply to all variables used in the study.

We use ERA5 daily surface air temperature at 2 metres averaged over a Western European region (SAT; latitudes: 45°N–56°N, longitudes: 5°W–16°E) as a proxy for near-surface thermal conditions, following Vautard et al. [2023]. SAT is used on the native spatial resolution of $0.25° \times 0.25°$ to compute the regional area-mean; this spatial mean time series is the target variable used in the RMM-VAE method throughout the study.

Large-scale atmospheric circulation is evaluated over the Euro-Atlantic sector using daily ERA5. The region of interest spans 30°N to 60°N latitude and 30°W to 20°E longitude, capturing key circulation dynamics affecting Western Europe. Compared to the classic North Atlantic-European domain commonly used in regime studies (e.g. Cattiaux et al. [2013]), this narrower domain improves cluster robustness and strengthens coupling to surface anomalies. Results for this larger domain are qualitatively similar and are reported in SI figure S1.

To compute weather regimes, we use the 500 hPa streamfunction (ψ), a common diagnostic for mid-tropospheric flow patterns, derived from ERA5 zonal and meridional wind components using the Climate Data Operators toolkit [Schulzweida, 2022]. In contrast to geopotential height, the streamfunction is not influenced by the background tropospheric warming due to anthropogenic climate change. This avoids the need for detrending and allows a more direct interpretation of circulation changes [Branstator, 2002]. The streamfunction ψ is computed in spherical coordinates as follows:

$$u_r = \frac{1}{r^2 \sin\theta}\frac{\partial \psi}{\partial \theta}, \qquad u_\theta = -\frac{1}{r \sin\theta}\frac{\partial \psi}{\partial r}, \tag{1}$$

where u_r and u_θ denote the meridional and zonal wind components, respectively, r is the radial distance from the Earth's centre, and θ is the zenith angle.

The streamfunction fields are regridded to a 2.5×2.5 spatial resolution to reduce data dimensionality and computational cost.

To analyse land–atmosphere coupling and hydrological influences associated with weather regimes, we further consider daily ERA5 soil moisture (SM) data in the uppermost layer (0–7 cm) and total precipitation (TP, mm), both regridded to the same $0.25° \times 0.25°$ grid as the SAT fields.

To remove the influence of the seasonal cycle and spatial variability, all fields are standardised at each grid point by subtracting the daily climatological mean and dividing by the daily standard deviation, computed over the full period. To reduce synoptic-scale variability and highlight more persistent circulation features, we apply a centred 5-day rolling mean to all daily fields. This smoothing approach suppresses transient weather disturbances while preserving the temporal resolution of the data [Rousi et al., 2021].

3.2 Definition of extremes and heatwaves

We define days of extreme heat, when the spatially averaged temperature in Western Europe exceeds the 90th percentile threshold computed over the full 1950–2022 period. This results in a

binary time series of extreme versus non-extreme summer days.

Beyond the simple definition based on the daily scalar target variable, we additionally classify Western European heatwaves by including thresholds for intensity, duration, and spatial extent Rousi et al. [2022], Rouges et al. [2023]. Here, heatwave days are defined when the following compound conditions are met:

- Intensity: Daily maximum temperature must exceed the 90th percentile, calculated using a 15-day centred moving window to account for seasonality.
- Duration: At least three consecutive days must exceed this local threshold.
- Spatial extent: On each of these days, at least 30% of the Western Europe domain must register grid point exceedances above their respective daily 90th percentile thresholds, computed from a centred 15-day climatology.

Finally, for quantifying inter-annual variability and long-term trends, we also define a seasonal index, reflecting the occurrence of high temperatures for each summer. This SAT_{90} index is the 90th percentile of the spatially averaged daily surface air temperature over Western Europe, computed separately for each summer.

3.3 Weather regime calculation

Two different methods are used for regime identification in this study: a conventional dimensionality reduction and clustering approach based on principal component analysis followed by k-means clustering (PCA and k-means), and a recently proposed targeted machine learning approach using a variational autoencoder trained on streamfunction fields (RMM-VAE).

3.3.1 PCA and k-means clustering As a baseline, we identify (non-targeted) weather regimes by applying PCA to daily JJA 500-hPa streamfunction anomalies over the Euro–Atlantic domain (Section 3.1), retaining the first ten PCs (∼96.4% variance) and then clustering their scores with k-means. For simplicity, we hereafter refer to regimes detected by this two-step approach simply as "k-means" regimes. Implementation and robustness details are provided in the SI.

3.3.2 RMM–VAE To complement the non-targeted baseline, we employ the targeted RMM–VAE approach introduced by Spuler et al. [2024]. This method learns a low-dimensional latent representation of the 500-hPa streamfunction and imposes a latent Gaussian-mixture prior to yield *probabilistic* daily weather regimes. A supervised term trains the encoder to predict the target variable, here averaged Western European near-surface temperature, thereby regularising the latent space toward patterns that are informative of the target variable while remaining physically consistent, i.e., constrained by dynamically plausible large-scale flow structures and variability modes rather than arbitrary statistical features (see also Spuler et al., 2024). Details about the method, architectural choices, training protocol, and hyperparameters are documented in the SI.

4 Results and discussion

4.1 Informativeness and statistical robustness of summer weather regimes

We compute summer weather regimes over the Euro-Atlantic region using both RMM-VAE and k-means. We first compare their performance across a range of choices for the cluster numbers ($K = 3$ to 10) by assessing three metrics: (1) the silhouette score describing their geometric separability [Rousseeuw, 1987], (2) the robustness to data sampling, and (3) the Brier skill score for predicting the target variable (temperatures in Western European) based on the weather regimes, to quantify regime informativeness (figure 1).

K-means generally shows a higher silhouette score (around 0.12 across most K) than RMM–VAE (peak ∼0.10 at $K = 3$, secondary ∼0.09 at $K = 5$; Fig. 1a), reflecting its optimisation of within-cluster variance. The absolute silhouette scores are modest, but consistent with previous studies on complex atmospheric regimes. For example, Happe et al. [2024] applied k-means and Gaussian mixture models to cluster spatiotemporal heatwaves and also observed low silhouette score across a range of K, attributed to the continuity of heatwave dynamics in phase space. This highlights that relatively low silhouette values may be characteristic of high-dimensional, continuous climate data, rather than a sign of poor clustering performance. The lower silhouette score for RMM-VAE is expected given its objective function, which prioritises latent structure learning and generative performance over explicit geometric compactness.

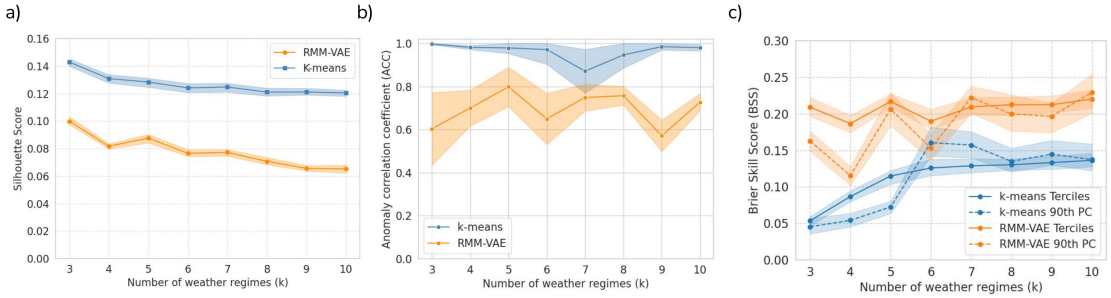

Figure 1. Comparison of performances of RMM-VAE (orange) and k-means clustering (blue) across different numbers of clusters $k = 3$ to 10. (a) Geometric separability, measured as the Silhouette scores. (b) Regime robustness, measured with the anomaly correlation coefficient (ACC), computed as the minimum accuracy of cluster label matches across 50 bootstrap samples (each using 90% of the data). (c) Informativeness of the clusters for the target variables, measured with the Brier skill score for predicting the below-, near-, and above-normal temperature terciles (solid lines) and the 90th percentile (dashed lines). In all plots, shading indicates the 95% confidence interval estimated via a non-parametric bootstrap of cases (N resampled with replacement, 1,000 replicates).

To assess regime stability under resampling, we randomly subsample 90% of the data 50 times and compare the resulting cluster labels with those from the full dataset. For each subsample, the minimum accuracy across all cluster matches is computed, providing a conservative stability measure. Consistent with the silhouette score, k-means is generally more robust (figure 1(b)). However, RMM-VAE achieves comparable robustness for $K = 5$ (mean accuracy ≈ 0.8), indicating stable, targeted regimes are attainable. Similar to our approach, Falkena et al. [2020] evaluated robustness via subsampling and also reported considerable sensitivity of regime patterns to data sampling and non-stationarity.

Finally, we evaluate the predictive skill of regime occurrence using the Brier skill score. This quantifies the added value of regime-based predictors over a climatological forecast for predicting daily temperatures. Here we assess both a tercile-based classification (below-, near-, and above-normal categories), and a binary 90th percentile threshold classification (SI Section 1). In essence, we ask: how much does knowing the weather regime tell us about the occurrence of different temperature states, including extreme conditions? RMM–VAE shows a clear advantage over k-means (Fig. 1c), achieving a Brier skill score roughly twice as large for the tercile-based classification across most K and maintaining superior skill for the 90th percentile threshold. Thus, despite slightly lower geometric compactness (Fig. 1a), RMM–VAE encodes features more predictive of summer SAT variability and extremes (Fig. 1c). This is consistent with findings that generative clustering can better capture latent atmospheric dynamics relevant for impact-relevant predictability [Spuler et al., 2024].

4.2 Physical relevance of summer weather regimes

Based on the combined evaluation of clustering robustness and silhouette and Brier skill scores, we select $K = 5$ as a suitable number of regimes for further analysis. To assess their physical relevance, figures 2-3 show the WRs derived from RMM-VAE and k-means, respectively, alongside composites of SAT, total precipitation (TP), and soil moisture (SM) anomalies.

RMM–VAE reveals clear circulation–surface links (Fig. 2). The warmest regime (WR1) exhibits a southerly-flow pattern with an anticyclonic anomaly over Central Europe and associated warm, dry advection; it features strong positive SAT anomalies, reduced TP, and SM deficits, and occurs on 21.7% of summer days (second most frequent). This WR1 closely matches the analogue-derived southerly pattern in Vautard et al. [2023]. The most frequent regime (WR5, 23.2%) is the coldest, with cyclonic anomalies over Central Europe and opposite surface impacts, effectively mirroring WR1. Upper-tropospheric meridional-wind composites across the Northern Hemisphere indicate that WR1 and WR5 form alternating ridges and troughs (SI Fig.S2), producing a wave-like, wavenumber-5 structure consistent with large-scale wave trains linked to mid-latitude extremes [Teng and Branstator, 2019]. The remaining regimes (WR2 17.7%, WR3 18.7%, WR4 18.6%) show weaker, more spatially variable SAT anomalies: WR2 features an Atlantic low extending toward Iberia, WR4 its counterpart with a North Atlantic high and weak cyclonic conditions over the Mediterranean, and WR3 is broadly anticyclonic with a ridge centred over the Mediterranean.

K-means captures broadly similar flow structures but with weaker 500-hPa signals and less

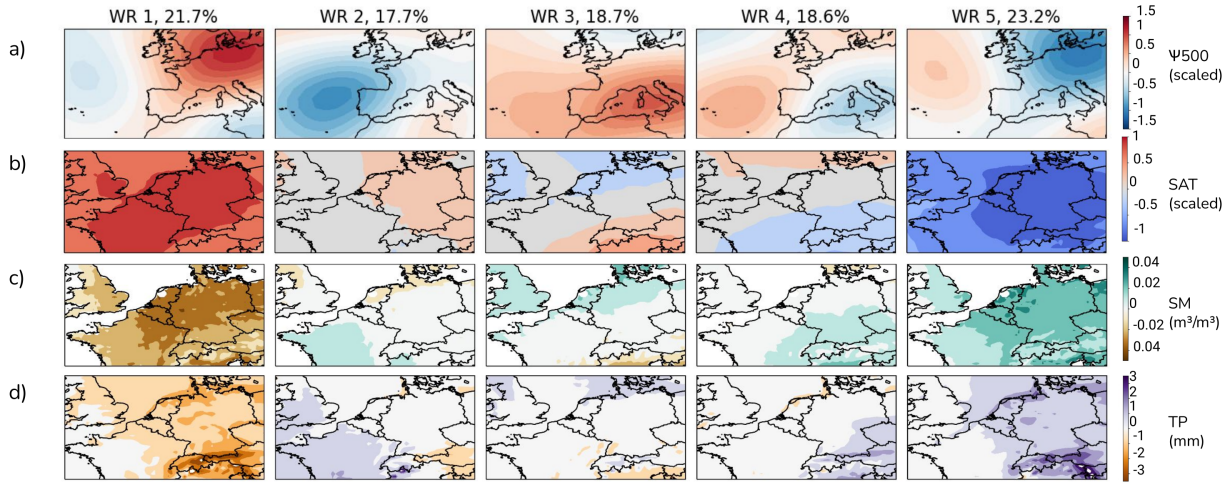

Figure 2. Weather regime composites from RMM-VAE clustering for $k = 5$. a) Streamfunction at 500 hPa showing the RMM-VAE regime patterns (cluster centres), with corresponding frequency of occurrence above each panel; b)–d) surface air temperature (SAT), soil moisture (SM), and total precipitation (TP) anomalies, respectively, composited with respect to each regime. The regimes are ordered in terms of the associated average temperature anomaly in the target region, with the first regime being associated with the highest temperatures.

coherent surface anomalies (Fig. 3). Its WR1 shows southerly flow, positive SAT anomalies, and SM deficits over Western Europe, but the mid-tropospheric anomaly is weaker and dry anomalies are more confined to the east. The most frequent k-means regime (WR3, 24.6%) features an anticyclone centred over the UK and yields coherent surface anomalies over the western subregion. Among the remaining regimes, WR2 (19.6%) is broadly anticyclonic with a weak Atlantic low northwest of the British Isles and warm anomalies in Western Europe, WR4 is the counterpart pattern with cool anomalies, and WR5 has a low centred over the North Sea with anticyclonic conditions across the southern domain.

These results demonstrate that compared to k-means, the RMM-VAE method identifies circulation patterns which better capture the SAT imprints in Western Europe, with additional associations in TP, and SM anomalies.

4.3 Linking summer weather regimes with heatwave occurrence

We next examine the co-occurrence of regimes and heat extremes in more detail. Southerly-flow regimes favour Western European heat extremes, whereas cyclonic configurations suppress temperatures and enhance precipitation [Vautard et al., 2023]. Figure 4(a) shows odds ratios of days with locally anomalous SAT exceeding the 90th percentile, given each of the WRs for RMM–VAE at the grid-point level. WR1 is strongly associated with extremes across most of the domain, with odds ratios exceeding 4, supporting its use as a predictive indicator, consistent with Fig1c. In contrast, k-means yields weaker and patchier associations (generally 2–3; Fig. 4c), indicating less regime specificity.

Figures 4(b,d) align WR occurrences with detected heatwaves, defined as ≥3 consecutive days with ≥30% of the domain above the local 90th percentile of daily SAT (Section 3.2). Over 1950–2022 we identify 59 events (330 days) with a median duration of 5 days and a maximum of 18 days (starting 2018-07-23, see SI Tables 1-2). RMM–VAE concentrates attribution with 92% of heatwave days coinciding with WR1 days (Fig. 4b). By contrast, k-means distributes heatwave days across WR1 (44%), WR2 (28%), and WR3 (25%)—the latter a UK-blocking pattern (Fig. 4d). This concentrated, physically interpretable linkage under RMM–VEA, thus, offers a better predictor for extremes, whereas k-means attribution dilutes predictive power.

4.4 Interannual variability of heat-linked weather regimes

To represent circulation variability on seasonal timescales, we use continuous WR indices defined as the projection of daily 500-hPa streamfunction anomalies onto each regime pattern [e.g. Mockert et al., 2024]. This yields continuous daily similarity time series for all k regimes; their JJA means are interpreted as continuous WR "frequencies," i.e. the seasonal dominance of each regime.

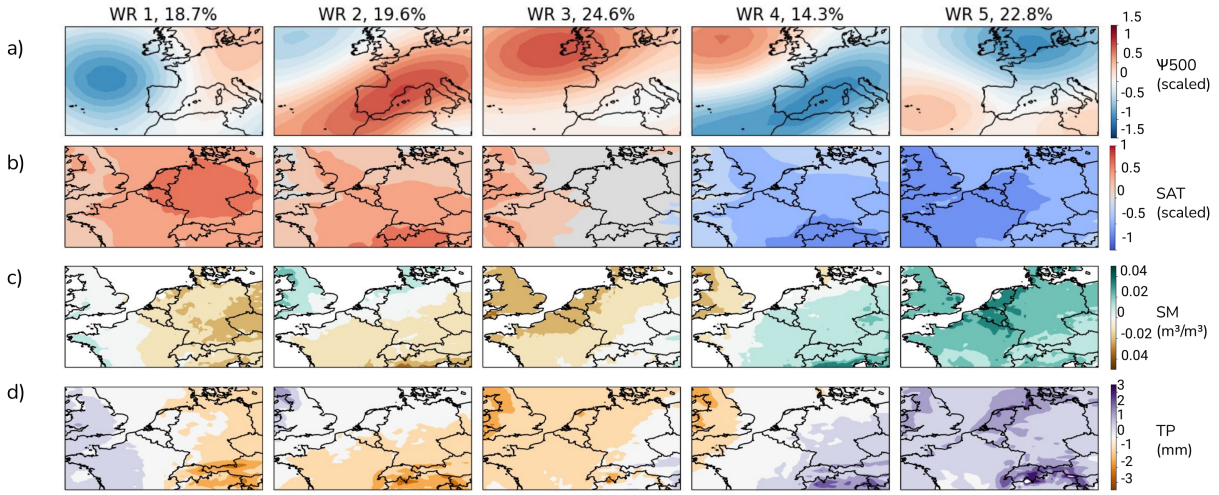

Figure 3. Same as Figure 2, but for weather regimes computed with k-means clustering for $k = 5$.

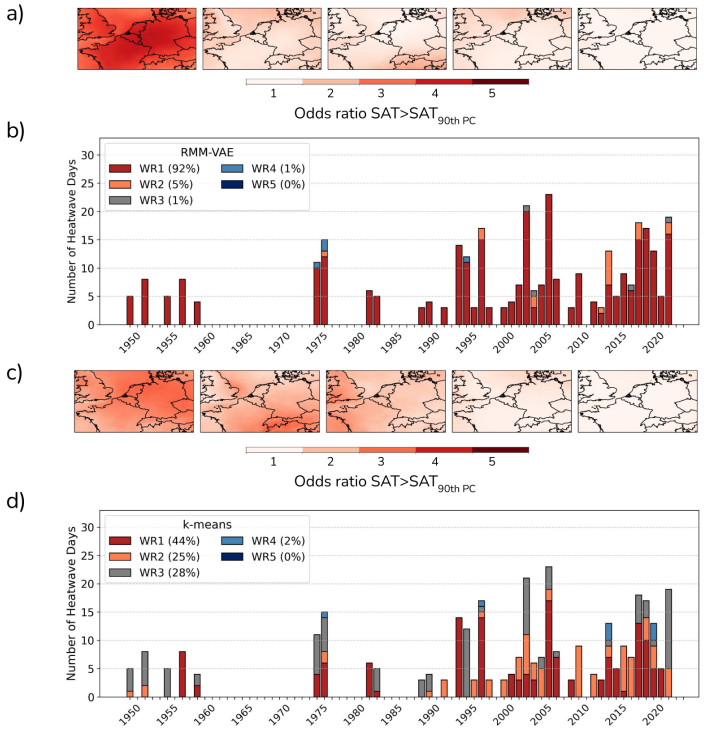

Figure 4. Heatwave occurrence and regime association over 1950–2022. a) Odds ratio of SAT exceeding the 90th percentile per grid cell in Western Europe, conditioned on each RMM-VAE regime; b) Number of heatwave days per summer identified based on the criterion of at least 3 consecutive days with SAT exceeding the 90th percentile over at least 30% of Western Europe, coloured by the corresponding RMM-VAE regime; c), d) same as a), b) but using the k-means regimes.

WR frequencies shifted markedly between 1950 and 2022: the heat-linked WR1 increases (+1.85% per decade in RMM–VAE; ∼+1.5% in k-means) while the cool-cyclonic WR5 declines (SI Fig. S4), with longer WR1 and shorter WR5 episodes (SI Fig. S5). These trends align with reports of more anticyclonic/warm-advection regimes over Europe contributing to rapid Western European summer warming [e.g. Horton et al., 2015, Rogers et al., 2022, Vautard et al., 2023].

The increase in WR1 and decline of WR5 seen in our analysis can be interpreted against a broader context of dynamical changes in the Euro-Atlantic summer circulation. Several lines of evidence indicate a weakened storm track activity, which favours more persistent, quasi-stationary flow patterns and thereby enhance the likelihood of prolonged warm anomalies over Europe [Coumou et al., 2015, Lehmann and Coumou, 2015]. Meridional shifts of the storm tracks, often linked to the summer NAO, further modulate blocking frequency and the distribution of hydroclimate anomalies across the basin [Dong et al., 2013]. In addition, more frequent double-jet configurations appear to provide a waveguide for amplified circumglobal wave patterns, sustaining warm-advection states over Western and Central Europe [Rousi et al., 2022]. These processes offer a physically consistent explanation for the observed strengthening of WR1 (and associated decrease in WR5 frequency). Both RMM-VAE and k-means capture these trends, though with differing magnitudes—the RMM-VAE highlights stronger circulation shifts, reflecting its closer alignment with surface heat impacts.

To quantify how interannual WR1 occurrence explains interannual SAT variability, we first regress out GMT from both series. This is analogous to linearly detrending both time series first, but appears more suitable because it directly controls for the (not exactly linear) common driver GMT, which here serves as our proxy for anthropogenic forcing. Regressing the residuals, that is, regressing the GMT-controlled SAT on GMT-controlled WR1 frequency is our estimate of how detrended seasonal WR1 frequency explains SAT variability. For RMM-VAE this yields $R^2 = 0.73$, indicating that 73% of Western European summer-mean temperature variability is explained by how often WR1 occurred that summer. Given the tight anti-correlation between WR1 and WR5 (r = -0.96), the latter adds little independent information. Thus, after accounting for GMT, how warm a summer is in the target region is largely set by how often a RMM-VAE WR1-type southerly flow occurs (or is absent). In contrast, for k-means the explained variance is only $R^2 = 0.42$, consistent with the less informative WR1 regime for extreme temperatures on weather timescales. The magnitude of our RMM-VAE fit compares well with Dong & Sutton's geopotential-based assessment of European summer variability based on detrended data ($R^2 \approx 0.61$; [Dong et al., 2013]). This indicates that a regime-frequency metric, designed to capture daily weather events, can capture a similar or even higher share of interannual SAT variance.

4.5 Explaining Western Europe temperature variability and trends

We next assess how summer SAT variability and long-term trends can be captured by WR1 occurrence and global mean warming within a single statistical–dynamical framework. To represent both interannual variability and forced trends, we use a simple multivariate linear regression model:

$$\mathrm{SAT} = a \cdot \mathrm{GMT} + b \cdot \mathrm{WR1} + \varepsilon, \tag{2}$$

where WR1 denotes the JJA mean of the continuous WR1 index and GMT is the global mean temperature anomaly. Here, a is interpreted as the *thermodynamic* contribution to Western European temperature variability (direct GMT→SAT), b as the *dynamic* contribution mediated by changes in the WR1 frequency. This regression thus provides a common framework for both explaining year-to-year fluctuations and decomposing long-term trends (see also Fig. S6 for a schematic of the assumed causal structure).

In Fig. 5 we evaluate regression skill for JJA SAT and SAT_{90} using WR1 frequency and GMT as predictors, for both RMM–VAE and k-means. Models are fitted over 1950–2022 (Fig. 5a,b) and with a 1950–2001 train / 2002–2023 test split (Fig. 5c,d); trained encoders/clusters are applied to the test years to assign regime labels under non-stationarity.

Over 1950–2022, RMM–VAE outperforms k-means: SAT $r = 0.92$ vs. 0.82; SAT_{90} 0.80 vs. 0.70 —all significant at 95%. Out of sample, RMM–VAE retains skill—SAT $r = 0.87$, SAT_{90} 0.70—whereas k-means degrades—SAT $r = 0.37$, SAT_{90} 0.36; the SAT correlation is not significant (Fig. 5c,d; SI Table 3). Consistent with these results, partial-R^2 maps (SI Fig. S7) show higher values for SAT than SAT_{90}, with maxima over France–Germany–Benelux.

For trends in the test period, RMM-VAE slightly underestimates SAT warming (0.35 $^\circ\mathrm{C\,dec^{-1}}$ vs. observed 0.44 $^\circ\mathrm{C\,dec^{-1}}$), while k-means shows a stronger underestimation (0.22 $^\circ\mathrm{C\,dec^{-1}}$). For SAT_{90}, RMM-VAE slightly overestimates the trend (0.42 vs. observed 0.38 $^\circ\mathrm{C\,dec^{-1}}$), whereas k-means underestimates it (0.28 $^\circ\mathrm{C\,dec^{-1}}$). The superior test-period performance of RMM-VAE

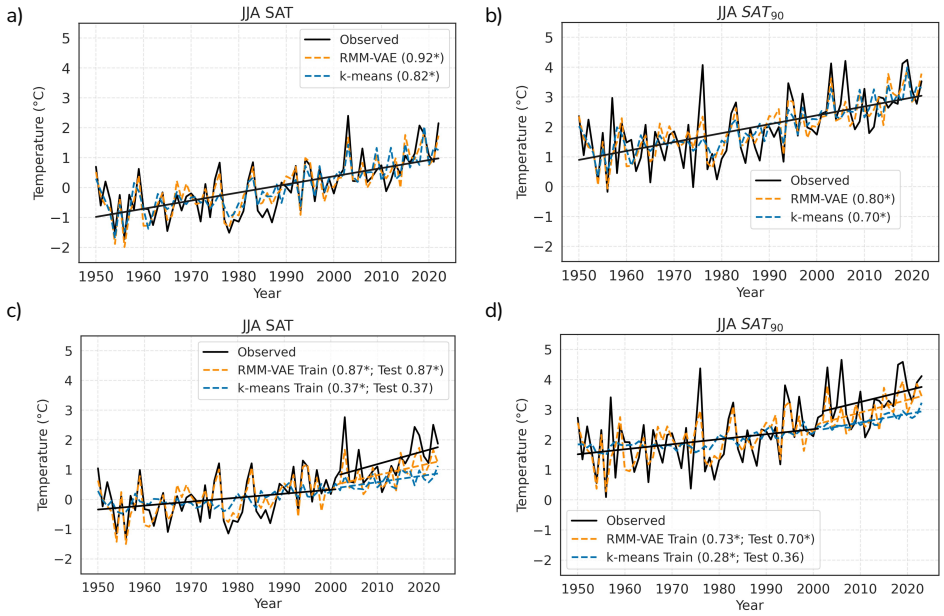

Figure 5. a) Time series of observed JJA SAT (black) in the target region for 1950-2022, alongside predicted SAT from linear regression models with seasonal weather regime 1 frequencies (based on RMM-VAE in orange, and based on k-means in blue) and GMT as predictors. Pearson correlation coefficients (r) are shown with a star if significant at the 95% level. The black line indicates the observed linear trend (trend values in SI Table 3). b) Same as a), but for predicting the 90th percentile of SAT (SAT_{90}). c) and d) same as a), and b) but using a train (1950–2001) and test (2002–2023) split. Trend lines for the test period are shown for observations, RMM-VAE, and k-means predictions.

indicates that it captures circulation modes more closely aligned with long-term warming trends, whereas k-means disperses trend-related variability across several clusters, limiting its generalisation.

Given this robust out-of-training performance, we now focus on RMM–VAE over the full period and use the same regression framework (Eq. 2) to disentangle the thermodynamic and dynamic components of the observed Western European temperature trend. Results for k-means are reported in the SI Table 4.

4.6 Disentangling the thermodynamic and dynamic components of observed temperature trends

Using the regression model introduced above (Eq.2), we quantify how both thermodynamic and dynamic processes contribute to observed SAT trends. To estimate the extent to which circulation itself responds to global warming, we additionally regress the WR1 on GMT:

$$\mathrm{WR1} = c\,\mathrm{GMT}, \tag{3}$$

so that c is our best estimate of the forced dynamical link (GMT → WR1). Note that internal variability also influences the frequency of WR1, which might affect the precision of this estimate.

Substituting the WR1–GMT relation into the SAT model gives:

$$\mathrm{SAT} = (a + bc)\,\mathrm{GMT}, \tag{4}$$

implying that the *total forced response* of SAT to GMT equals $a + bc$, thus the sum along the direct thermodynamic and indirect (via dynamics) pathway (see also Kretschmer et al. [2021], Cosford et al. [2025]). The remaining variability reflects internal circulation changes not linearly tied to GMT. Here we get estimates of $a = 1.33$, $b = 4.96$, and $c = 0.10$.

Over the period from 1950 to 2022, GMT warmed by 0.14 ± 0.01 °C per decade, while Western Europe warmed approximately twice as much (0.27 ± 0.04 °C per decade). This yields an "excess" regional warming of 0.13 ± 0.04 °C per decade, defined as the difference between Western European and GMT (Table 1). During this same period, WR1 had a trend of 0.02 ± 0.01 per decade (i.e., a +1.85% per decade increase, SI Fig. S4), indicating more frequent occurrences of WR1-type southerly flow.

Plugging in these values into our linear regression model (Eq.2), we get an estimate of 0.28 ± 0.03 °C per decade, closely matching the observed Western European warming

Table 1. Observed and predicted temperature trends for Western European summer (JJA) based on the seasonal RMM–VAE WR1 index. *SAT* is the observed Western European mean temperature trend, and *GMT* the global mean SAT trend (indicated in °C per decade) . The *Excess* is the regional warming beyond the global mean (*SAT–GMT*). The *Thermodynamic* contribution reflects the direct GMT–SAT link, the *Dynamical* contribution reflects the observed circulation-induced SAT change, and the *Forced dynamical* term represents the GMT-driven WR1 contribution. Percentage rows report the shares of SAT and Excess. Uncertainty is shown as *estimate* ± half-width of the 95% confidence interval: trends use a non-parametric bootstrap of year–value pairs (1,000 replicates); fractions use first-order (delta) propagation of uncertainty for ratios $z = y/x$ with uncorrelated numerator and denominator, $u_z \approx |z|\sqrt{(u_y/y)^2 + (u_x/x)^2}$, with CIs given by $z \pm 1.96\, u_z$ (treating regression coefficients as fixed).

Observed trends (°C dec^{-1})	
SAT (Western Europe)	0.27 ± 0.04
GMT (Global mean)	0.14 ± 0.01
Excess (SAT–GMT)	0.13 ± 0.04
RMM–VAE estimations (trend ± CI, °C dec^{-1})	
Thermodynamic (a·GMT)	0.19 ± 0.01
Dynamical (b·WR1)	0.09 ± 0.03
Forced dynamical (b·c·GMT)	0.07 ± 0.00
Fractional contributions (%)	
Dynamical / SAT	$34\% \pm 13\%$
Forced dynamical / SAT	$25\% \pm 4\%$
Thermodynamic / SAT	$69\% \pm 11\%$
Dynamical / Excess	$70\% \pm 32\%$
Forced dynamical / Excess	$52\% \pm 15\%$
Thermodynamic / Excess	$35\% \pm 10\%$

(0.27 ± 0.04 °C per decade). Thus, this simple model can explain all the observed SAT warming trend in the target region. The direct, thermodynamic-induced warming trend, here represented by $a \cdot \mathrm{GMT}$, is 0.19 ± 0.01 °C per decade, accounting for about $69 \pm 11\%$ of the total SAT trend, which aligns with the estimate of $65.5\% \pm 3.2\%$ for Eurasia reported by Hu et al. [2025]. The dynamic contribution, $b \cdot \dot{\mathrm{W}}\mathrm{R1}$, amounts to 0.09 ± 0.03 °C per decade ($34 \pm 13\%$).

Using the regression coefficient c from model Eq.3, we estimate the forced component of the WR1 trend to be 0.01 ± 0.00 °C per decade. This implies a forced dynamical contribution to the SAT trend of $b \cdot c \cdot \dot{\mathrm{G}}\mathrm{MT} = 0.07 \pm 0.00$ °C per decade, corresponding to $25 \pm 4\%$ of the total SAT trend. Substituting this forced circulation response into the SAT model (Eq.4), the combined thermodynamic and forced dynamical effects, $(a + bc)$, gives a total forced response for Western European SAT of 0.25 ± 0.01 °C per decade. This explains about $94 \pm 15\%$ of the total observed SAT trend, with the remaining residual trend (0.02 ± 0.04 °C per decade) reflecting internal circulation variability not linearly tied to GMT.

Our estimate of 34% for the dynamical contribution aligns with previous findings, including 24% reported by Vautard et al. [2023] for 1950–2022 and the 33% found by Singh et al. [2023] for 1979–2021, both of which analysed, however, trends in hot extremes ($\mathrm{T_{xx}}$) rather than mean SAT. Therefore, we apply the same regression framework (Eqs.2-4) for the summer 90th-percentile SAT ($\mathrm{SAT_{90}}$), yielding a dynamical contribution to the observed trend of $29\% \pm 11\%$, in close agreement with those studies (SI Table 5). Compared with Dong and Sutton [2025], who used a summer geopotential height-based index over a broad European box, we obtain larger circulation shares for Western Europe: $34\% \pm 13\%$ of the total and $70\% \pm 32\%$ of the excess, versus $19\% \pm 6\%$ and $29\% \pm 10\%$. This likely reflects our subregional focus and the use of a WR1 frequency predictor from RMM-VAE rather than a spatial-mean Z_{500} index.

In summary, our results confirm that GMT is the dominant driver of mean and extreme summer SAT [Hu et al., 2025, Dong and Sutton, 2025, Singh et al., 2023, Vautard et al., 2023], but that circulation changes make a non-negligible contribution to excess European warming, consistent with studies highlighting more frequent southerly flow and blocking [Rousi et al., 2022, Van Oldenborgh et al., 2022].

5 Summary and conclusions

This study applies the targeted probabilistic clustering approach RMM-VAE [Spuler et al., 2024] to study drivers of Western European temperature extremes and their trends. It compares RMM-VAE with the standard approach of combining PCA and k-means to characterise summer circulation patterns in Western Europe. While k-means performs better on geometric separation and sampling robustness, RMM-VAE produces physically coherent regimes more directly linked to regional temperature impacts. In this way, the method shifts the focus from purely statistical compactness toward impact-relevance, offering a framework better suited to disentangling the circulation processes that shape European summer heat. Previous applications of RMM-VAE include identifying circulation patterns relevant to extreme winter precipitation in Morocco [Spuler et al., 2024, 2025]. Its probabilistic formulation assigns likelihoods to each circulation state, capturing transitional states often overlooked in hard clustering [Falkena et al., 2023] and providing a more complete picture of circulation variability relevant for predictability and impact attribution.

Our results demonstrate three main advances. First, RMM-VAE yields physically consistent regimes, with a single southerly-flow state (WR1) concentrating the vast majority of heatwave days, in contrast to the more diffuse attribution across k-means clusters. Second, combining WR1 frequency with GMT in a linear regression setup, provides a skilful and interpretable model to predict interannual variability and trends in both mean ($R^2 = 84\%$ of explained trend) and extreme summer SAT ($R^2 = 65\%$ of explained trend), with comparable predictive skill persisting in out-of-sample tests (that is, when the regimes and hence the regression model are trained on a subsample only). This simple linear regression model provides a compact representation of the circulation–temperature relationship, integrating WR1 frequency and GMT into a form directly compatible with statistical–dynamical forecasting frameworks that use targeted predictors to refine ensemble-based forecasts [Dobrynin et al., 2018, Carvalho-Oliveira et al., 2022, Famooss Paolini et al., 2025]. Third, attribution analysis based on the same linear regression model shows that 34% of the observed Western European warming trends is explained by dynamic changes (the increase in WR1 frequency), which we estimate here to be partly forced, while acknowledging remaining uncertainty regarding the role of internal variability, consistent with recent attribution studies of dynamical contributions to European heat [Rousi et al., 2022, Vautard et al., 2023, Dong and Sutton, 2025]. In terms of excess European warming, observed WR1-circulation changes explain 70% of the total summer warming beyond the global mean signal.

Our results confirm the central role of southerly flow in driving Western European summer warming, in line with analogue-based studies that emphasise this pattern as a key contributor to heat extremes [Vautard et al., 2023]. What distinguishes our approach is that the RMM-VAE identifies this circulation state within a full regime classification, rather than isolating a single index or pattern a priori. This ensures that while the dominant warm-advection regime (WR1) emerges clearly as the main driver, other dynamically coherent regimes are simultaneously represented, providing a more complete picture of the circulation spectrum. In contrast, single-pattern approaches maximise informativeness but risk overlooking variability associated with other regimes, whereas standard WR classifications based on k-means capture the full phase space but diffuse heat impacts across multiple clusters. By combining targeted learning with a regime framework, the RMM-VAE strikes a balance between impact-relevance and dynamical representativeness, yielding a more robust basis for attribution.

Future work should test the robustness of targeted regime approaches across domains, variables, and models, and explore their integration with statistical–dynamical postprocessing. Extending the analysis beyond Western Europe could assess circumglobal circulation trends [Teng and Branstator, 2019], while accounting for decadal internal variability, which may complicate the detection of regional trends. Furthermore, applying the method to seasonal forecasts, for example, could clarify its added predictive value, while testing it on climate model ensembles could reveal whether models that better capture observed circulation trends also better reproduce Europe's excess warming—raising the possibility of emergent constraints on future projections [Hall et al., 2019].

Beyond attribution, the RMM-VAE regimes provide compact and interpretable predictors with direct relevance for prediction and risk assessment. Their ability to concentrate heat extremes into a single circulation driver, while still representing transitional states probabilistically, makes them valuable for both statistical–dynamical postprocessing and climate services. In this context, further assessing the role of local modulating factors relevant for heat extremes, such as vegetation and soil moisture conditions is important. Moreover, extending the approach to other impact-relevant variables such as precipitation or compound events could help bridge the gap between dynamical understanding, prediction skill, and long-term projections.

Funding
The authors acknowledge funding from the Horizon Europe project EXPECT (Towards an Integrated Capability to Explain and Predict Regional Climate Changes), Grant Agreement No. 101137656. F.S. further acknowledges funding from the Advancing the Frontiers of Earth System Prediction (AFESP) Doctoral Training Programme at the University of Reading.

Author contributions
J.O. and M.K. designed the research; J.O. performed the research; J.O., F.S. and M.K. contributed to the methodology; J.O. analysed and curated data; J.O. and F.S. prepared the software and visualisation; J.O., F.S. and M.K. contributed to the interpretation of results and editing of the manuscript; J.O. led the writing.

References
Siyi Chen, Yufei Liu, Yuanyuan Yi, Yiling Zheng, Jun Yang, Tiantian Li, Ta-Chien Chan, Rui Duan, Shenjing He, and Cui Guo. Long-term impacts of heatwaves on accelerated ageing. *Nature Climate Change*, pages 1–8, 2025.

Sharon Campbell, Tomas A Remenyi, Christopher J White, and Fay H Johnston. Heatwave and health impact research: A global review. *Health & place*, 53:210–218, 2018.

Corey Lesk, Pedram Rowhani, and Navin Ramankutty. Influence of extreme weather disasters on global crop production. *Nature*, 529(7584):84–87, 2016.

T Armada Brás, SG Simoes, Filipa Amorim, and Patricia Fortes. How much extreme weather events have affected european power generation in the past three decades? *Renewable and Sustainable Energy Reviews*, 183:113494, 2023.

Jannis Von Buttlar, Jakob Zscheischler, Anja Rammig, Sebastian Sippel, Markus Reichstein, Alexander Knohl, Martin Jung, Olaf Menzer, M Altaf Arain, Nina Buchmann, et al. Impacts of droughts and extreme-temperature events on gross primary production and ecosystem respiration: a systematic assessment across ecosystems and climate zones. *Biogeosciences*, 15(5): 1293–1318, 2018.

Efi Rousi, Andreas H Fink, Lauren S Andersen, Florian N Becker, Goratz Beobide-Arsuaga, Marcus Breil, Giacomo Cozzi, Jens Heinke, Lisa Jach, Deborah Niermann, et al. The extremely hot and dry 2018 summer in central and northern europe from a multi-faceted weather and climate perspective. *Natural Hazards and Earth System Sciences*, 23(5):1699–1718, 2023.

Xinhui Li, Jiayu Zheng, Chunzai Wang, Xiayan Lin, and Zhixiong Yao. Unraveling the roles of jet streams on the unprecedented hot july in western europe in 2022. *npj Climate and Atmospheric Science*, 7(1):323, 2024.

Efi Rousi, Kai Kornhuber, Goratz Beobide-Arsuaga, Fei Luo, and Dim Coumou. Accelerated western european heatwave trends linked to more-persistent double jets over eurasia. *Nature communications*, 13(1):3851, 2022.

Jitendra Singh, Sebastian Sippel, and Erich M Fischer. Circulation dampened heat extremes intensification over the midwest usa and amplified over western europe. *Communications Earth & Environment*, 4(1):432, 2023.

Buwen Dong and Rowan T Sutton. Drivers and mechanisms contributing to excess warming in europe during recent decades. *npj Climate and Atmospheric Science*, 8(1):41, 2025.

Peter Pfleiderer, Anna Merrifield, István Dunkl, Homer Durand, Enora Cariou, Julien Cattiaux, and Sebastian Sippel. The contribution of circulation changes to summer temperature trends in the northern hemisphere mid-latitudes: A multi-method quantification. *EGUsphere*, 2025:1–28, 2025.

Daniel E Horton, Nathaniel C Johnson, Deepti Singh, Daniel L Swain, Bala Rajaratnam, and Noah S Diffenbaugh. Contribution of changes in atmospheric circulation patterns to extreme temperature trends. *Nature*, 522(7557):465–469, 2015.

Dim Coumou, Jascha Lehmann, and Johanna Beckmann. The weakening summer circulation in the northern hemisphere mid-latitudes. *Science*, 348(6232):324–327, 2015.

Jascha Lehmann and Dim Coumou. The influence of mid-latitude storm tracks on hot, cold, dry and wet extremes. *Scientific reports*, 5(1):17491, 2015.

Dim Coumou, Vladimir Petoukhov, Stefan Rahmstorf, Stefan Petri, and Hans Joachim Schellnhuber. Quasi-resonant circulation regimes and hemispheric synchronization of extreme weather in boreal summer. *Proceedings of the National Academy of Sciences*, 111(34): 12331–12336, 2014.

Kai Kornhuber, Vladimir Petoukhov, Stefan Petri, Stefan Rahmstorf, and Dim Coumou. Evidence for wave resonance as a key mechanism for generating high-amplitude quasi-stationary waves in boreal summer. *Climate Dynamics*, 49(5):1961–1979, 2017.

Geert Jan Van Oldenborgh, Michael F Wehner, Robert Vautard, Friederike EL Otto, Sonia I Seneviratne, Peter A Stott, Gabriele C Hegerl, Sjoukje Y Philip, and Sarah F Kew. Attributing and projecting heatwaves is hard: we can do better. *Earth's Future*, 10(6):e2021EF002271, 2022.

Robert Vautard, Julien Cattiaux, Tamara Happé, Jitendra Singh, Rémy Bonnet, Christophe Cassou, Dim Coumou, Fabio D'andrea, Davide Faranda, Erich Fischer, et al. Heat extremes in western europe increasing faster than simulated due to atmospheric circulation trends. *Nature Communications*, 14(1):6803, 2023.

Clara Deser, Adam Phillips, Vincent Bourdette, and Haiyan Teng. Uncertainty in climate change projections: the role of internal variability. *Climate dynamics*, 38(3):527–546, 2012.

Clara Deser and Adam S Phillips. A range of outcomes: the combined effects of internal variability and anthropogenic forcing on regional climate trends over europe. *Nonlinear Processes in Geophysics*, 30(1):63–84, 2023.

Rong Zhang, Rowan Sutton, Gokhan Danabasoglu, Young-Oh Kwon, Robert Marsh, Stephen G Yeager, Daniel E Amrhein, and Christopher M Little. A review of the role of the atlantic meridional overturning circulation in atlantic multidecadal variability and associated climate impacts. *Reviews of Geophysics*, 57(2):316–375, 2019.

Sebastian Sippel, Nicolai Meinshausen, Anna Merrifield, Flavio Lehner, Angeline G Pendergrass, Erich Fischer, and Reto Knutti. Uncovering the forced climate response from a single ensemble member using statistical learning. *Journal of Climate*, 32(17):5677–5699, 2019.

Haiyan Teng, Ruby Leung, Grant Branstator, Jian Lu, and Qinghua Ding. Warming pattern over the northern hemisphere midlatitudes in boreal summer 1979–2020. *Journal of Climate*, 35(11): 3479–3494, 2022.

SE Perkins-Kirkpatrick and SC Lewis. Increasing trends in regional heatwaves. *Nature communications*, 11(1):3357, 2020.

Elizaveta Felsche, Andrea Böhnisch, Benjamin Poschlod, and Ralf Ludwig. European hot and dry summers are projected to become more frequent and expand northwards. *Communications Earth & Environment*, 5(1):410, 2024.

Emmanuel Rouges, Laura Ferranti, Holger Kantz, and Florian Pappenberger. European heatwaves: Link to large-scale circulation patterns and intraseasonal drivers. *International Journal of Climatology*, 43(7):3189–3209, 2023.

Tamara Happe, JS Wijnands, Miguel Ángel Fernández-Torres, Paolo Scussolini, Laura Muntjewerf, and Dim Coumou. Detecting spatiotemporal dynamics of western european heatwaves using deep learning. *Artificial Intelligence for the Earth Systems*, 3(4):e230107, 2024.

Christophe Cassou, Laurent Terray, and Adam S Phillips. Tropical atlantic influence on european heat waves. *Journal of climate*, 18(15):2805–2811, 2005.

Sarah E Perkins. A review on the scientific understanding of heatwaves—their measurement, driving mechanisms, and changes at the global scale. *Atmospheric Research*, 164:242–267, 2015.

Kai Kornhuber, Dim Coumou, Elisabeth Vogel, Corey Lesk, Jonathan F Donges, Jascha Lehmann, and Radley M Horton. Amplified rossby waves enhance risk of concurrent heatwaves in major breadbasket regions. *Nature Climate Change*, 10(1):48–53, 2020.

Giorgia Di Capua, S Sparrow, Kai Kornhuber, Eftychia Rousi, S Osprey, D Wallom, B Van Den Hurk, and Dim Coumou. Drivers behind the summer 2010 wave train leading to russian heatwave and pakistan flooding. *npj Climate and Atmospheric Science*, 4(1):55, 2021.

Daniela IV Domeisen, Elfatih AB Eltahir, Erich M Fischer, Reto Knutti, Sarah E Perkins-Kirkpatrick, Christoph Schär, Sonia I Seneviratne, Antje Weisheimer, and Heini Wernli. Prediction and projection of heatwaves. *Nature Reviews Earth & Environment*, 4(1):36–50, 2023.

Ming Luo, Ngar-Cheung Lau, and Zhen Liu. Different mechanisms for daytime, nighttime, and compound heatwaves in southern china. *Weather and Climate Extremes*, 36:100449, 2022.

Philipp Zschenderlein, Andreas H Fink, Stephan Pfahl, and Heini Wernli. Processes determining heat waves across different european climates. *Quarterly Journal of the Royal Meteorological Society*, 145(724):2973–2989, 2019.

Sonia I Seneviratne, Thierry Corti, Edouard L Davin, Martin Hirschi, Eric B Jaeger, Irene Lehner, Boris Orlowsky, and Adriaan J Teuling. Investigating soil moisture–climate interactions in a changing climate: A review. *Earth-Science Reviews*, 99(3-4):125–161, 2010.

Almudena García-García, Francisco José Cuesta-Valero, Diego G Miralles, Miguel D Mahecha, Johannes Quaas, Markus Reichstein, Jakob Zscheischler, and Jian Peng. Soil heat extremes can outpace air temperature extremes. *Nature Climate Change*, 13(11):1237–1241, 2023.

Julianna Carvalho-Oliveira, Giorgia Di Capua, Leonard F Borchert, Reik V Donner, and Johanna Baehr. Causal relationships and predictability of the summer east atlantic teleconnection. *Weather and Climate Dynamics*, 5(4):1561–1578, 2024.

Erica Madonna, Camille Li, Christian M Grams, and Tim Woollings. The link between eddy-driven jet variability and weather regimes in the north atlantic-european sector. *Quarterly Journal of the Royal Meteorological Society*, 143(708):2960–2972, 2017.

Alexander Lemburg and Andreas H Fink. Investigating the medium-range predictability of european heatwave onsets in relation to weather regimes using ensemble reforecasts. *Quarterly Journal of the Royal Meteorological Society*, 150(764):3957–3988, 2024.

Paul-Antoine Michelangeli, Robert Vautard, and Bernard Legras. Weather regimes: Recurrence and quasi stationarity. *Journal of Atmospheric Sciences*, 52(8):1237–1256, 1995.

Christian M Grams, Remo Beerli, Stefan Pfenninger, Iain Staffell, and Heini Wernli. Balancing europe's wind-power output through spatial deployment informed by weather regimes. *Nature climate change*, 7(8):557–562, 2017.

Abdel Hannachi, David M Straus, Christian LE Franzke, Susanna Corti, and Tim Woollings. Low-frequency nonlinearity and regime behavior in the northern hemisphere extratropical atmosphere. *Reviews of Geophysics*, 55(1):199–234, 2017.

Julianna Carvalho-Oliveira, Leonard F Borchert, Eduardo Zorita, and Johanna Baehr. Self-organizing maps identify windows of opportunity for seasonal european summer predictions. *Frontiers in Climate*, 4:844634, 2022.

Takuya Kurihana, Ilijana Mastilovic, Lijing Wang, Aurelien Meray, Satyarth Praveen, Zexuan Xu, Milad Memarzadeh, Alexander Lavin, and Haruko Wainwright. Identifying climate patterns using clustering autoencoder techniques. *Artificial Intelligence for the Earth Systems*, 3(3): e230035, 2024.

Elizaveta Felsche, Andrea Böhnisch, and Ralf Ludwig. Inter-seasonal connection of typical european heatwave patterns to soil moisture. *npj Climate and Atmospheric Science*, 6(1):1, 2023.

Hannah C Bloomfield, David J Brayshaw, and Andrew J Charlton-Perez. Characterizing the winter meteorological drivers of the european electricity system using targeted circulation types. *Meteorological Applications*, 27(1):e1858, 2020.

Fiona R Spuler, Marlene Kretschmer, Yevgeniya Kovalchuk, Magdalena Alonso Balmaseda, and Theodore G Shepherd. Identifying probabilistic weather regimes targeted to a local-scale impact variable. *Environmental Data Science*, 3:e25, 2024.

Fiona R Spuler, Marlene Kretschmer, Magdalena Alonso Balmaseda, Yevgeniya Kovalchuk, and Theodore G Shepherd. Learning predictable and informative dynamical drivers of extreme precipitation using variational autoencoders. *Weather and Climate Dynamics*, 6(3):995–1014, 2025.

Hans Hersbach, Bill Bell, Paul Berrisford, Shoji Hirahara, András Horányi, Joaquín Muñoz-Sabater, Julien Nicolas, Carole Peubey, Raluca Radu, Dinand Schepers, et al. The era5 global reanalysis. *Quarterly journal of the royal meteorological society*, 146(730):1999–2049, 2020.

Julien Cattiaux, Benjamin Quesada, Ara Arakélian, Francis Codron, Robert Vautard, and Pascal Yiou. North-atlantic dynamics and european temperature extremes in the ipsl model: sensitivity to atmospheric resolution. *Climate dynamics*, 40(9):2293–2310, 2013.

Uwe Schulzweida. Cdo user guide (2.1. 0). zenodo, 2022.

Grant Branstator. Circumglobal teleconnections, the jet stream waveguide, and the north atlantic oscillation. *Journal of Climate*, 15(14):1893–1910, 2002.

Eftychia Rousi, Frank Selten, Stefan Rahmstorf, and Dim Coumou. Changes in north atlantic atmospheric circulation in a warmer climate favor winter flooding and summer drought over europe. *Journal of Climate*, 34(6):2277–2295, 2021.

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Swinda KJ Falkena, Jana de Wiljes, Antje Weisheimer, and Theodore G Shepherd. Revisiting the identification of wintertime atmospheric circulation regimes in the euro-atlantic sector. *Quarterly Journal of the Royal Meteorological Society*, 146(731):2801–2814, 2020.

Haiyan Teng and Grant Branstator. Amplification of waveguide teleconnections in the boreal summer. *Current Climate Change Reports*, 5(4):421–432, 2019.

Fabian Mockert, Christian M Grams, Sebastian Lerch, Marisol Osman, and Julian Quinting. Multivariate post-processing of probabilistic sub-seasonal weather regime forecasts. *Quarterly Journal of the Royal Meteorological Society*, 150(765):4771–4787, 2024.

Cassandra DW Rogers, Kai Kornhuber, Sarah E Perkins-Kirkpatrick, Paul C Loikith, and Deepti Singh. Sixfold increase in historical northern hemisphere concurrent large heatwaves driven by warming and changing atmospheric circulations. *Journal of Climate*, 35(3):1063–1078, 2022.

Buwen Dong, Rowan T Sutton, Tim Woollings, and Kevin Hodges. Variability of the north atlantic summer storm track: mechanisms and impacts on european climate. *Environmental Research Letters*, 8(3):034037, 2013.

Marlene Kretschmer, Samantha V Adams, Alberto Arribas, Rachel Prudden, Niall Robinson, Elena Saggioro, and Theodore G Shepherd. Quantifying causal pathways of teleconnections. *Bulletin of the American Meteorological Society*, 102(12):E2247–E2263, 2021.

Lara R Cosford, Rohit Ghosh, Marlene Kretschmer, Clare Oatley, and Theodore G Shepherd. Estimating the contribution of arctic sea-ice loss to central asia temperature anomalies: the case of winter 2020/2021. *Environmental Research Letters*, 20(3):034007, 2025.

Yuhan Hu, Zhiyan Zuo, Haishan Chen, and Wenjian Hua. Attribution of multidecadal summer temperature variations over eurasia. *Environmental Research Letters*, 20(5):054005, 2025.

Swinda KJ Falkena, Jana de Wiljes, Antje Weisheimer, and Theodore G Shepherd. A bayesian approach to atmospheric circulation regime assignment. *Journal of Climate*, 36(24):8619–8636, 2023.

Mikhail Dobrynin, Daniela IV Domeisen, Wolfgang A Müller, Louisa Bell, Sebastian Brune, Felix Bunzel, André Düsterhus, Kristina Fröhlich, Holger Pohlmann, and Johanna Baehr. Improved teleconnection-based dynamical seasonal predictions of boreal winter. *Geophysical Research Letters*, 45(8):3605–3614, 2018.

Luca Famooss Paolini, Paolo Ruggieri, Salvatore Pascale, Erika Brattich, and Silvana Di Sabatino. Hybrid statistical–dynamical seasonal prediction of summer extreme temperatures in europe. *Quarterly Journal of the Royal Meteorological Society*, 151(766):e4900, 2025.

Alex Hall, Peter Cox, Chris Huntingford, and Stephen Klein. Progressing emergent constraints on future climate change. *Nature Climate Change*, 9(4):269–278, 2019.

Supplementary Information for the paper "Targeted weather regimes identify circulation patterns behind Western European summer heat extremes and trends"

Julianna Carvalho-Oliveira[1,*], Fiona R. Spuler [1,2] and Marlene Kretschmer[1]
[1]Leipzig Institute for Meteorology, Leipzig University, Leipzig, Germany
[2]Department of Meteorology, University of Reading, Reading, UK
*Author to whom any correspondence should be addressed.
Contact: julianna.carvalho_oliveira@uni-leipzig.de

1 Skill Scores

1.1 Brier Skill Score

To evaluate the predictive skill of regime occurrence, we compute Brier Skill Scores (BSS) for classifying daily surface air temperature (SAT) into terciles and for identifying extremes above the 90th percentile over Western Europe. Regime occurrence frequencies are used as predictors in a probabilistic classification framework.
The Brier Score (BS) is a strictly proper scoring rule that quantifies the mean squared difference between predicted probabilities and the actual outcomes. It is defined as:

$$\mathrm{BS} = \frac{1}{N}\sum_{n=1}^{N}\sum_{j=1}^{m}\left(\delta_{nj} - p_j\right)^2, \tag{1}$$

where N is the number of time steps, m is the number of forecast categories, p_j is the forecast probability for category j, and δ_{nj} is the Kronecker delta (1 if the observation at time n is in category j, 0 otherwise).
To assess the added value of the regime-based forecasts, we compute the BSS relative to a climatological forecast:

$$\mathrm{BSS} = 1 - \frac{\mathrm{BS}_{\mathrm{forecast}}}{\mathrm{BS}_{\mathrm{climatology}}}, \tag{2}$$

where $\mathrm{BS}_{\mathrm{forecast}}$ is the BS of the forecast based on regime occurrence and $\mathrm{BS}_{\mathrm{climatology}}$ is the BS assuming climatological probabilities. Higher BSS values indicate improved predictive skill relative to climatology.

1.2 Silhuette Score

The geometric separability of clusters can be measured in terms of the silhouette score [Rousseeuw, 1987]. This score combines measures of cohesion and separation. Cohesion reflects the average distance between a data point and other members of the same cluster, while separation quantifies the average distance between a point and members of the nearest-neighbouring cluster. The silhouette score ranges from –1 to 1, with higher values indicating more distinct and well-separated clusters. A score near zero suggests overlapping clusters, and negative values imply incorrect clustering assignments.

2 Weather Regime Computation

2.1 PCA and K-means Clustering

We apply PCA, commonly referred to as empirical orthogonal function (EOF) analysis in atmospheric sciences, to the preprocessed 500 hPa streamfunction daily anomalies during boreal summer (JJA) over the Euro-Atlantic domain as described in Manuscript Section 3.1.
The first ten PCs are retained, explaining approximately 96.4% of the total variance in the streamfunction field. These leading modes represent the dominant patterns of mid-tropospheric variability. Next, k-means clustering is applied to the associated PC time series to identify recurrent circulation patterns. To ensure stable and robust regime definitions, the clustering is repeated with multiple initialisations and a fixed random seed. The PCA is implemented using the `Eof` class from the `eofs` Python package, and k-means clustering is carried out using `KMeans` from `scikit-learn`. The resulting PCA+k-means clusters define the weather regimes used for subsequent analysis, and hereafter will be referred to simply as k-means clusters.

2.2 Regression Mixture Model Variational Autoencoder (RMM-VAE)

To identify targeted weather regimes, we apply the RMM-VAE method introduced by Spuler et al. [2024]. The RMM-VAE integrates dimensionality reduction and probabilistic clustering within a unified statistical model, improving upon conventional approaches that typically separate these steps and rely on hard cluster assignments.
The first key innovation of this method is the use of a Gaussian mixture model in the latent space (that is, the dimensionality reduced space), allowing the identification of weather regimes as probabilistic clusters rather than assuming a single multivariate Gaussian distribution. This enables a softer cluster assignment for each day, better capturing transitional states between regimes.
The second enhancement involves incorporating a supervised learning component into the VAE structure: the encoder is trained to predict a target variable—in our case, near-surface temperature over Western Europe. This prediction serves to regularise the latent space (i.e. adds a constraint during training that favours simple, target-consistent representations), promoting cluster structures that reflect coherent responses in the target variable. As shown in Spuler et al. [2024], this regularisation disentangles latent dimensions associated with variations in the target and results in weather regimes with more consistent impacts on surface weather.
These methodological advances have been achieved through an implemented modified loss function that balances reconstruction accuracy, latent space regularisation, and prediction performance. Details of the methodology, including derivations of the loss function are described in Spuler et al. [2024].
Here, the RMM-VAE is implemented as a fully connected neural network with three dense layers of decreasing dimensionality (128, 64, and 32 units), using ReLU activation functions throughout. A latent space of dimension 10 is used to capture the main circulation features while allowing sufficient complexity for regime separation. Training is performed using a batch size of 128 and a fixed learning rate of 0.001 over 100 epochs. To account for uncertainty due to random weight initialisation, the model is trained 50 times with different initial seeds. In each run, the model is first trained on a train–test split, and the final validation loss is recorded. The model instance with the lowest validation loss is selected, and its weights are used to retrain the architecture on the full dataset. This two-step approach ensures robust model selection and consistency across initialisations. All neural network components are implemented using the `Keras` Python library [Chollet et al., 2015].

References

Peter J Rousseeuw. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65, 1987.

Fiona R Spuler, Marlene Kretschmer, Yevgeniya Kovalchuk, Magdalena Alonso Balmaseda, and Theodore G Shepherd. Identifying probabilistic weather regimes targeted to a local-scale impact variable. *Environmental Data Science*, 3:e25, 2024.

François Chollet et al. keras, 2015.

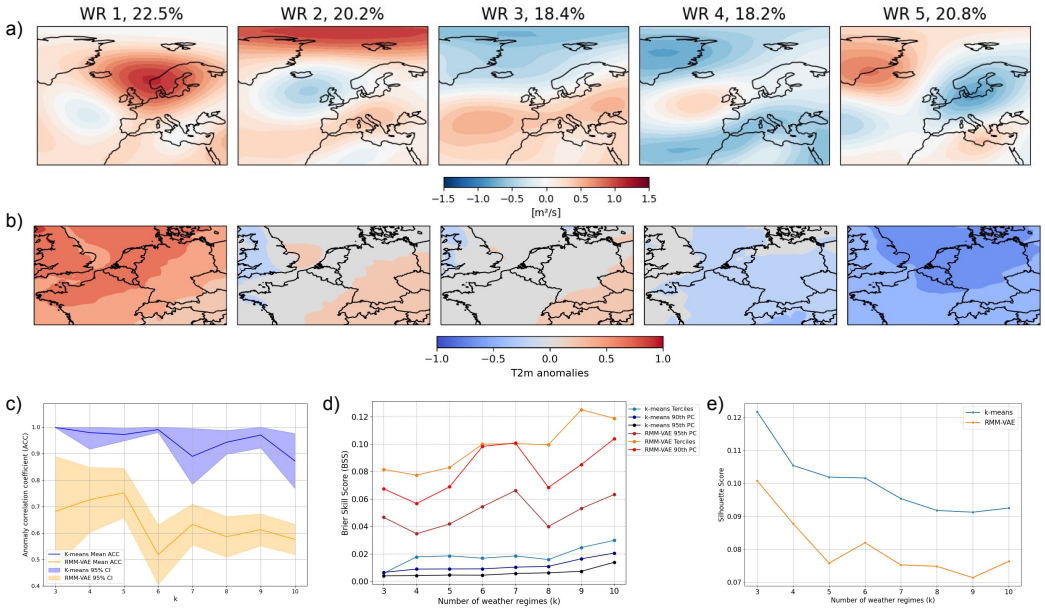

Figure 1: (a) Weather regimes for the North Atlantic–European (NAE) sector at $k = 5$ derived with RMM-VAE from 500 hPa streamfunction (sf500). (b) Composites of surface air temperature (SAT) anomalies over Western Europe associated with the regimes in (a). (c) Regime robustness quantified by a conservative anomaly correlation coefficient (ACC): the minimum cluster label-match accuracy across 50 bootstrap resamples (each using 90% of the data), shown for RMM-VAE (yellow) and k-means (blue). (d) Brier Skill Score (BSS) for predicting temperature terciles and the 90th/95th percentiles, evaluated for RMM-VAE (tones of red) and k-means (tones of blue) across $k = 3$–10. (e) Silhouette scores for RMM-VAE (orange) and k-means (blue) across $k = 3$–10.

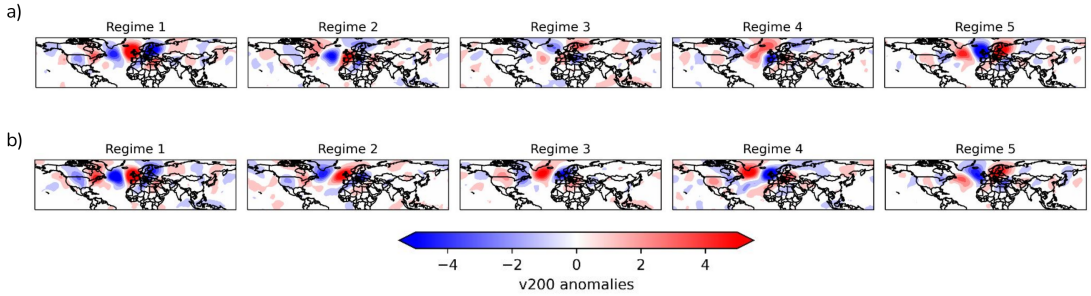

Figure 2: Composites of meridional wind anomalies at 200 hPa (v200) for each summer weather regime: (a) RMM-VAE and (b) k-means.

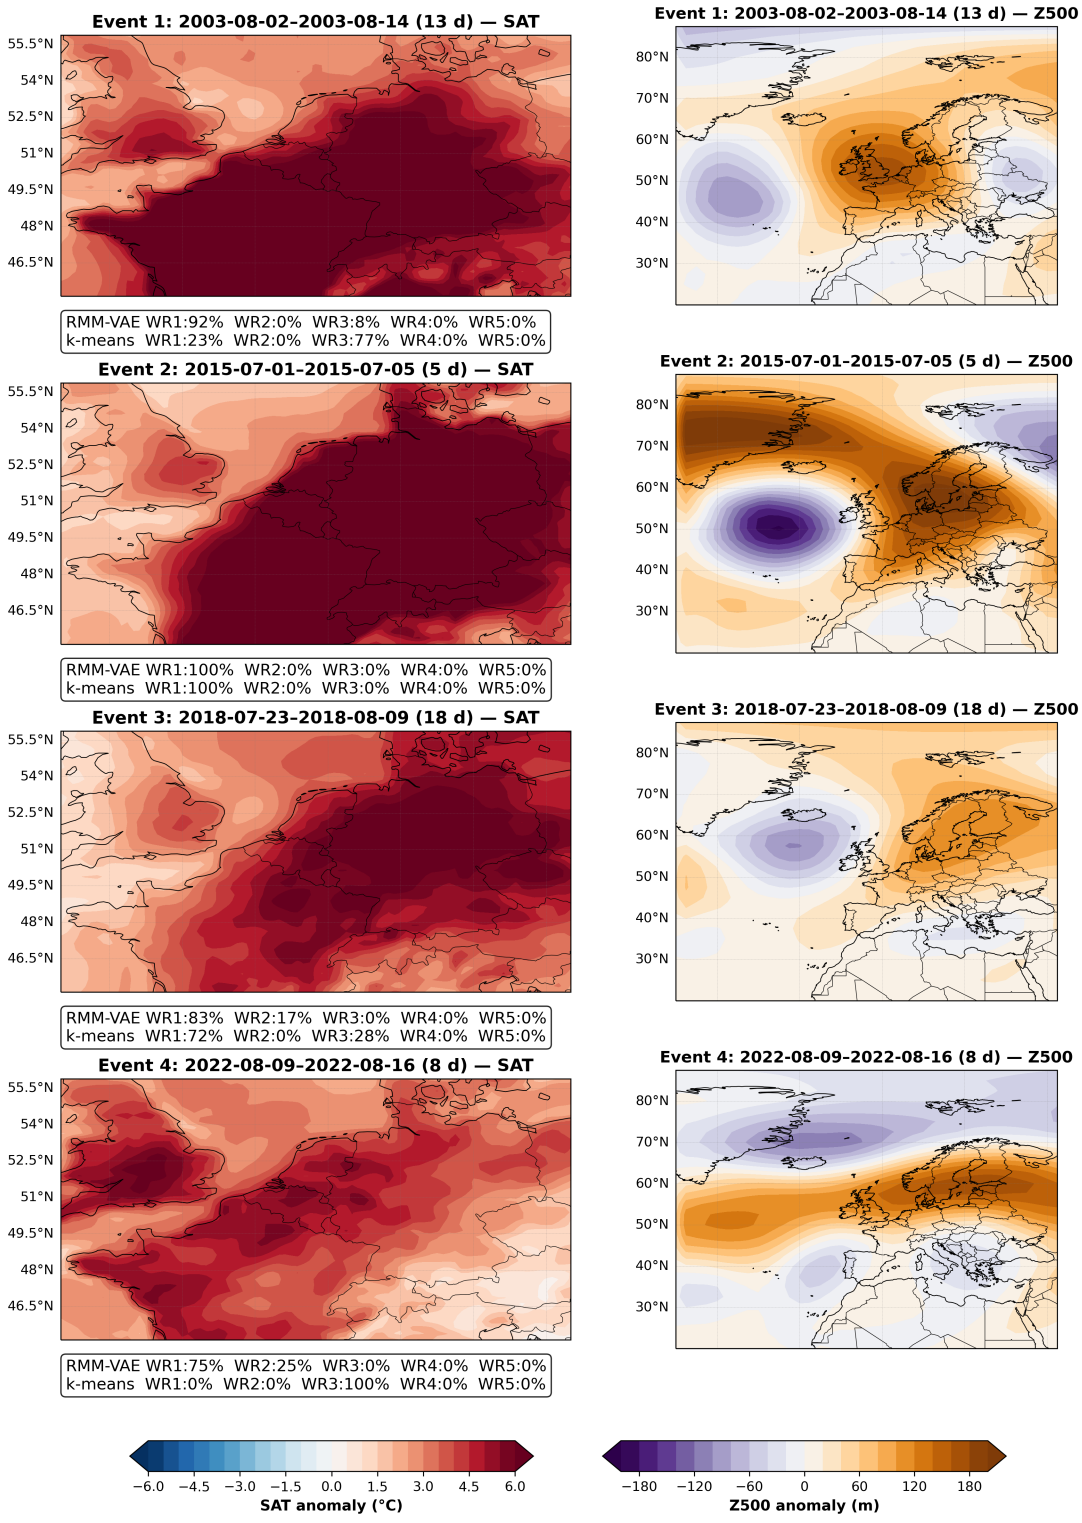

Figure 3: Examples of four major Western European heatwaves in 2003, 2015, 2018, and 2022. For each event, the left panel shows the surface air temperature (SAT) anomaly map, with the event date and duration indicated in the title. Below each SAT map, bars indicate the percentage of days classified into regimes WR1–WR5 by RMM-VAE and k-means. The right panel displays the corresponding 500 hPa geopotential height (Z500) anomaly pattern.

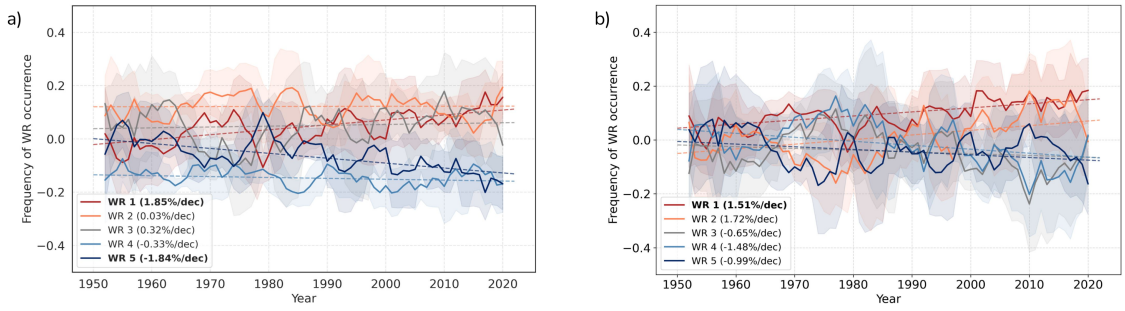

Figure 4: Trends in seasonal frequency of occurrence for each of the five weather regimes from a) RMM-VAE and b) k-means over the period 1950–2022, including linear trend lines. Trends in bold are significant at the 95% level according to a two-sided Student´s t-test.

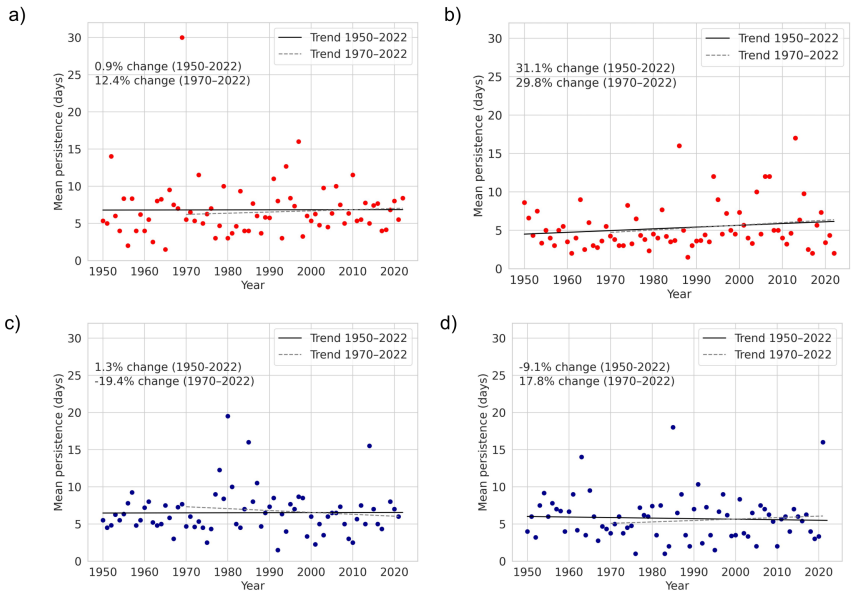

Figure 5: Time series of annual mean persistence (in days) of WR1 (a-b) and WR5 (c-d) for the period 1950–2022. The black solid line shows the linear trend over the full period (1950–2022), while the dashed line indicates the trend for 1970–2022. The latter is included to account for the strong outlier in 1969 for a), which substantially affects the full-period estimate. Percentage change in persistence per decade is indicated for each trend. (a-c) RMM-VAE; (b-d) k-means.

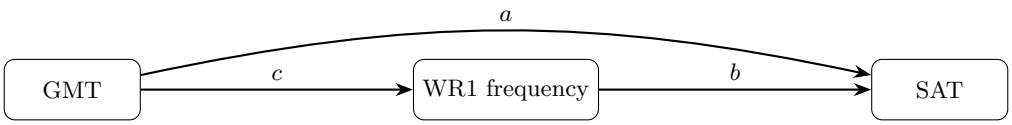

Total effect of GMT on SAT: $a + b\,c$

Figure 6: Directed acyclic graph (DAG) illustrating the decomposition of contributions: direct thermodynamic pathway a (GMT→SAT), forced dynamical pathway c (GMT→WR1) combined with dynamic sensitivity b (WR1→SAT), yielding a total effect $a + bc$.

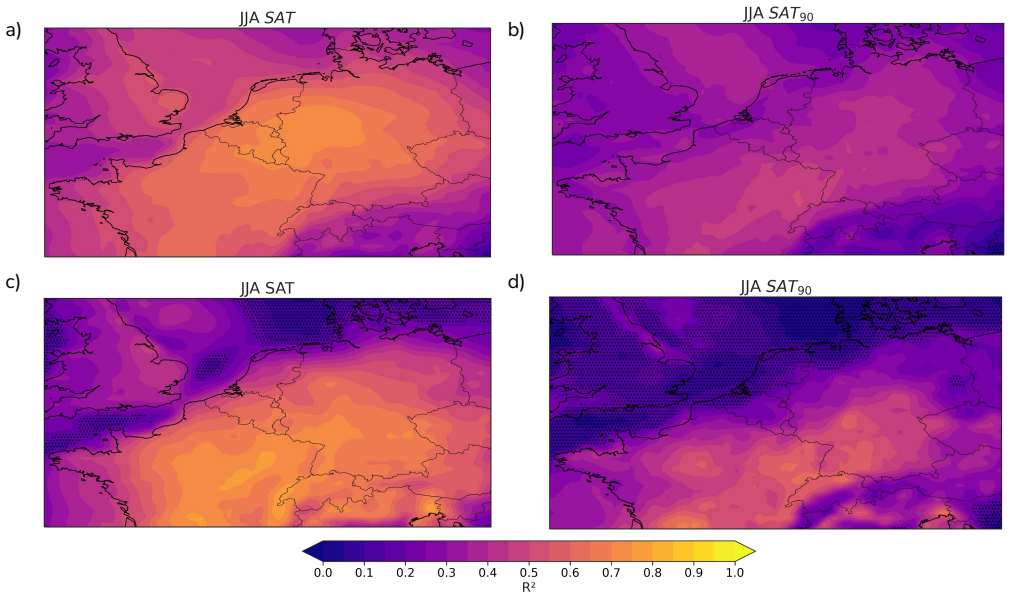

Figure 7: Maps of partial R^2 for Western Europe showing the variance in summer temperature explained by WR1 frequency conditional on GMT for results from RMM-VAE. (a) SAT for 1950-2022; (b) SAT_{90} for 1950-2022; (c) SAT for the independent test period 2002-2023; (d) SAT_{90} for 2002-2023. Partial R^2 at each grid point is computed as the squared correlation between GMT-residualised SAT and GMT-residualised WR1. Stippling marks statistical insignificant grid points with $p >= 0.05$; pixels with insufficient data are masked.

Event	Start	End	Duration (days)
HW01	1950-06-04	1950-06-08	5
HW02	1952-06-29	1952-07-06	8
HW03	1955-08-20	1955-08-24	5
HW04	1957-06-29	1957-07-01	3
HW05	1957-07-03	1957-07-07	5
HW06	1959-07-08	1959-07-11	4
HW07	1975-08-02	1975-08-12	11
HW08	1976-06-23	1976-07-07	15
HW09	1982-06-01	1982-06-06	6
HW10	1983-07-09	1983-07-13	5
HW11	1989-06-18	1989-06-20	3
HW12	1990-08-01	1990-08-04	4
HW13	1992-08-07	1992-08-09	3
HW14	1994-07-22	1994-07-28	7
HW15	1994-07-30	1994-08-01	3
HW16	1994-08-03	1994-08-06	4
HW17	1995-07-30	1995-08-03	5
HW18	1995-08-17	1995-08-23	7
HW19	1996-06-06	1996-06-08	3
HW20	1997-08-08	1997-08-14	7
HW21	1997-08-16	1997-08-25	10
HW22	1998-08-09	1998-08-11	3
HW23	2000-06-19	2000-06-21	3
HW24	2001-08-23	2001-08-26	4
HW25	2002-06-16	2002-06-19	4
HW26	2002-08-17	2002-08-19	3
HW27	2003-06-03	2003-06-05	3
HW28	2003-06-10	2003-06-14	5
HW29	2003-08-02	2003-08-14	13

Table 1: Detected Western Europe heatwave spells (1950–2022; part A). Events are defined as at least three consecutive days with at least 30% domain exceedance of the local daily 90th percentile (15-day centred climatology).

Event	Start	End	Duration (days)
HW30	2004-06-08	2004-06-10	3
HW31	2004-08-08	2004-08-10	3
HW32	2005-06-18	2005-06-24	7
HW33	2006-06-10	2006-06-13	4
HW34	2006-07-01	2006-07-06	6
HW35	2006-07-16	2006-07-28	13
HW36	2007-06-07	2007-06-11	5
HW37	2007-07-14	2007-07-16	3
HW38	2009-06-30	2009-07-02	3
HW39	2010-07-08	2010-07-16	9
HW40	2012-08-18	2012-08-21	4
HW41	2013-07-26	2013-07-28	3
HW42	2014-06-07	2014-06-12	6
HW43	2014-07-18	2014-07-20	3
HW44	2014-07-23	2014-07-26	4
HW45	2015-07-01	2015-07-05	5
HW46	2016-07-18	2016-07-20	3
HW47	2016-08-23	2016-08-28	6
HW48	2017-06-17	2017-06-23	7
HW49	2018-07-23	2018-08-09	18
HW50	2019-06-24	2019-06-29	6
HW51	2019-07-22	2019-07-26	5
HW52	2019-08-24	2019-08-29	6
HW53	2020-06-23	2020-06-26	4
HW54	2020-08-06	2020-08-14	9
HW55	2021-06-16	2021-06-20	5
HW56	2022-06-15	2022-06-18	4
HW57	2022-07-17	2022-07-20	4
HW58	2022-08-09	2022-08-16	8
HW59	2022-08-23	2022-08-25	3

Table 2: Continuation of Part A; see Table 1 for definition.

Table 3: Summary of linear regression model performance using WR1 frequency and GMT as predictors of summer mean SAT and 90th percentile SAT (SAT_{90}) over Western Europe. The table complements Fig.5 in the manuscript, providing numerical values for predicted trends, Pearson correlation coefficients (r), and coefficients of determination (R^2) for both the RMM-VAE and k-means models. Results are shown for the full period ("in-sample", 1950–2022), the training period (1950–2001), and the independent test period (2002–2023). An asterisk indicates significance of the Pearson correlation at the 95% level. The *Observed* column reports reference SAT and SAT_{90} trends only (correlation and R^2 are not applicable).

	Observed			**RMM-VAE**			**k-means**		
	In-sample	**Train**	**Test**	**In-sample**	**Train**	**Test**	**In-sample**	**Train**	**Test**
Trend SAT	0.27	0.13	0.44	0.28	0.16	0.35	0.27	0.10	0.22
r (p) SAT	–	–	–	0.92*	0.87*	0.87*	0.82*	0.37*	0.37
R² SAT	–	–	–	0.84	0.76	0.40	0.67	0.14	-0.77
Trend SAT90	0.30	0.18	0.35	0.30	0.20	0.44	0.29	0.13	0.30
r (p) SAT90	–	–	–	0.80*	0.73*	0.70*	0.70*	0.28*	0.36
R² SAT90	–	–	–	0.65	0.54	0.27	0.50	0.08	-0.78

Table 4: Observed and decomposed JJA SAT trends over Western Europe (°C per decade) based on the k-means WR1 index. *SAT* is the observed Western European mean temperature trend, and *GMT* the global mean SAT trend. The *Excess* is the regional warming beyond the global mean (*SAT–GMT*). The *Thermodynamic* contribution reflects the direct GMT–SAT link; the *Dynamical* contribution reflects the observed circulation-induced SAT change; and the *Forced dynamical* term represents the GMT-driven WR1 response. Percentage rows report the shares of SAT and Excess. Uncertainty is shown as *estimate* ± half-width of the 95% confidence interval: trends use a non-parametric bootstrap of year–value pairs (1,000 replicates); fractions use first-order (delta) propagation of uncertainty for ratios $z = y/x$ with uncorrelated numerator and denominator, $u_z \approx |z|\sqrt{(u_y/y)^2 + (u_x/x)^2}$, with CIs given by $z \pm 1.96\, u_z$ (treating regression coefficients as fixed).

Observed trends (°C dec^{-1})	
SAT (Western Europe)	0.27 ± 0.04
GMT (Global mean)	0.14 ± 0.01
Excess (SAT–GMT)	0.13 ± 0.04
k-means estimations (trend ± CI, °C dec^{-1})	
Thermodynamic (a·GMT)	0.21 ± 0.01
Dynamical (b·WR1)	0.06 ± 0.03
Forced dynamical (b·c·GMT)	0.04 ± 0.00
Fractional contributions (%)	
Dynamical / SAT	$22\% \pm 10\%$
Forced dynamical / SAT	$15\% \pm 2\%$
Thermodynamic / SAT	$79\% \pm 13\%$
Dynamical / Excess	$46\% \pm 24\%$
Forced dynamical / Excess	$31\% \pm 9\%$
Thermodynamic / Excess	$55\% \pm 16\%$

Table 5: Same as Table 4, but for JJA 90th-percentile SAT (SAT_{90}) over Western Europe (°C per decade). SAT_{90} denotes the Western European SAT_{90} trend. Results are shown for WR1 indices derived from RMM–VAE and k-means.

	RMM–VAE	**k-means**
Observed trends (°C dec^{-1})		
SAT_{90} (Western Europe)	0.30 ± 0.05	0.30 ± 0.05
GMT (Global mean)	0.14 ± 0.01	0.14 ± 0.01
Excess (SAT_{90}–GMT)	0.16 ± 0.04	0.16 ± 0.04
Model-based decompositions (trend ± CI, °C dec^{-1})		
Thermodynamic (a·GMT)	0.21 ± 0.01	0.24 ± 0.01
Dynamical (b·WR1)	0.09 ± 0.03	0.05 ± 0.02
Forced dynamical (b·c·GMT)	0.06 ± 0.00	0.03 ± 0.00
Fractional contributions (%)		
Dynamical / SAT_{90}	$29\% \pm 11\%$	$17\% \pm 8\%$
Forced dynamical / SAT_{90}	$22\% \pm 4\%$	$12\% \pm 2\%$
Thermodynamic / SAT_{90}	$72\% \pm 12\%$	$82\% \pm 14\%$
Dynamical / Excess	$56\% \pm 25\%$	$32\% \pm 17\%$
Forced dynamical / Excess	$41\% \pm 12\%$	$22\% \pm 6\%$