Diagnosing concurrent drivers of weather extremes: application to hot and cold days in North America

Davide Faranda, Gabriele Messori, Pascal Yiou

Diagnosing concurrent drivers of weather extremes: application to warm and cold days in North America

Davide Faranda · Gabriele Messori · Pascal Yiou

Received: date / Accepted: date

Abstract A fundamental challenge in climate science is decomposing the concurrent drivers of weather extremes in observations. Achieving this can provide insights into the drivers of individual extreme events as well as into possible future changes in extreme event frequencies under greenhouse forcing. In the present work, we exploit recent results from dynamical systems theory to study the co-variation and recurrence statistics of different atmospheric fields. Specifically, we present a methodology to quantify the recurrences of bivariate fields and the coupling between distinct univariate fields in terms of their joint recurrences. The coupling is defined by a parameter which varies according to the chosen fields, season, and domain and can be understood in terms of the underlying physics of the atmosphere. For suitably chosen fields, this approach enables to decompose the different drivers of weather extremes. Here, we compute the above metrics for near-surface temperature and sea level pressure, and use them to study warm or cold days over North America. We first identify states where temperature extremes are strongly or weakly coupled to the large-scale atmospheric circulation, and then elucidate the interplay between coupling and the occurrence of temperature extremes.

Keywords Weather extremes · Analogues · Climate Dynamics · Dynamical Systems Theory

D. Faranda
Laboratoire des Sciences du Climat et de l'Environnement LSCE-IPSL, CEA Saclay l'Orme des Merisiers, UMR 8212 CEA-CNRS-UVSQ, Université Paris-Saclay, 91191 Gif-sur-Yvette, France & London Mathematical Laboratory, 8 Margravine Gardens, London, W68RH, UK.
E-mail: davide.faranda@lsce.ipsl.fr

G. Messori
Department of Earth Sciences, Uppsala University, Uppsala, Sweden & Department of Meteorology and Bolin Centre for Climate Research, Stockholm University, Stockholm, Sweden

P. Yiou
Laboratoire des Sciences du Climat et de l'Environnement LSCE-IPSL, CEA Saclay l'Orme des Merisiers, UMR 8212 CEA-CNRS-UVSQ, Université Paris-Saclay, 91191 Gif-sur-Yvette, France

1 Introduction

Weather extremes are complex events, driven by a variety of physical processes. Disentangling these processes is a fundamental challenge in climate science and is essential in order to reconstruct the chain leading to individual or co-occurring extremes (Davies, 2015; Messori et al, 2016; Harnik et al, 2016; Kornhuber et al, 2019). Early statistical studies of weather extremes adopted a univariate framework, projecting extreme events onto peak-over-threshold (Pickands III, 1975) or block-maxima (Gnedenko, 1943) of single observables (e.g. storm surges, wind speeds or temperatures (Galambos et al, 1994)). Multivariate approaches were introduced to take into account combinations of variables that achieve simultaneously large or small values (e.g. wind-speed and storm surge). Recently, the awareness has grown that individual variables may not be extreme themselves, but that their joint occurrence may yield an extremal behavior (see e.g. Bevacqua et al (2017)).

In order to understand the root drivers of large-scale weather extremes, one must therefore follow all the relevant variables in space and time over extended regions. For example, heatwaves or cold days reflect the interaction between persistent circulation regimes and temperature patterns. In some cases additional confounding factors, such as soil-moisture levels, further complicate the picture (Zscheischler and Seneviratne, 2017; Hirschi et al, 2011). Studying these interplays a priori necessitates a large number of numerical experiments with multi-parameter models.

In this paper, we aim to combine a multivariate view of extremes with analytical and computational efficiency, by constraining the analyses with the generic behavior of chaotic systems. We propose a new methodology, based on the adaptation and extension of recent mathematical results from dynamical systems theory, and apply it to the study of temperature extremes (warm or cold days) over North America. Our work builds upon the numerous studies that have tried to connect the dynamical properties of the atmosphere to the occurrence of temperature extremes in the region (e.g. Grotjahn (2016) and references therein, Messori et al (2016); Harnik et al (2016)), as well as the broader literature looking at mid-latitude extremes (e.g. Von Storch and Zwiers (2001); Coumou et al (2014); Palmer (2013); Gálfi et al (2017); Messori et al (2019)). These point to the degree of coupling between atmospheric circulation variables and temperature as a key step to both understand the physical drivers of the extremes and constrain the variability of climate models in order to better simulate such events in future climates. Here, we specifically diagnose the dynamical features and coupling of the sea-level pressure (SLP) and 2-m temperature (T2M) fields in reference to the above-mentioned North American temperature extremes. The purpose is to provide a proof-of-concept for the applicability of our novel analysis approach – which is entirely general and may in theory be extended to any number of variables – to the study of multi-variate atmospheric configurations.

We begin by providing a theoretical definition of three dynamical systems metrics: d, θ^{-1} and α. The local dimension d can be intuitively interpreted

as a proxy for the number of active degrees of freedom of the system around a given state. The persistence θ^{-1} measures the residence time around such state. Finally, the coupling parameter (also referred to as co-recurrence ratio) α informs on the dependence structure between instantaneous configurations of different variables. The underlying concept is akin to that of joint recurrence plots (Marwan et al, 2007). We next test these metrics on simple stochastic processes and discrete dynamical systems. Following this, we apply our framework to study the drivers underlying the occurrence of temperature extremes in North America.

2 Dynamical systems indicators

2.1 A dynamical systems framework

The attractor of a dynamical system is a geometrical object defined in the space hosting all the possible states of the system (phase space). For atmospheric flows, it is impossible to obtain all *variables* of the system, as this would require the knowledge of the properties of each fluid parcel. What is instead available is a set, or sequences, of *observables* – namely transformations of the variables of the system. By making some assumptions on the dimensionality of the system, it is possible to retrieve its entire dynamics from the observables (Huke, 2006). This allows the computation of the dimension of the system (Grassberger, 1983; Grassberger and Procaccia, 1984). The major caveat of such approach lies in the underlying mathematical assumptions, that render the application to geophysical datasets unreliable (Eckmann and Ruelle, 1992). Recently, an alternative approach to determine a system's dynamical properties has been proposed, which does not require a priori assumptions on the dimensionality of the system. This approach is based on an analysis of the recurrences of trajectories of the dynamical system, and has been successfully tested on various complex datasets (Faranda et al, 2017b,a; Messori et al, 2017; Faranda et al, 2019a). One of the outcomes is a characterization of each point ζ along the sequence of observables of the system by two dynamical indicators: the local dimension (d) and the persistence (θ^{-1}) (Faranda et al, 2017b). There have been theoretical (Barros et al, 2019) and experimental (Faranda et al, 2017a,c) arguments that show that recurrences of observables are useful to investigate properties of the underlying phase space of atmospheric motions, even though all the variables defining the phase space of the system cannot be taken into account for obvious computational reasons.

Before detailing the computation of d and θ^{-1}, we give a physical interpretation to these seemingly abstract dynamical indicators. The local dimension d measures the geometry of the trajectories in a small region of the phase space. It is therefore linked to the number of degrees of freedom that a system can locally explore. In other terms, a region of the phase space with a low (high) dimension, allows a small (large) number of different configurations of the system. In Faranda et al (2017b) and Messori et al (2017), we showed

that lower dimensional states of the atmospheric circulation over the North Atlantic correspond to a predominantly zonal flow. Higher local dimensions are instead found for blocked flows. The persistence θ^{-1} informs on the persistence of the system in the same small region of the phase space. A higher (lower) persistence implies that trajectories will evolve more slowly (rapidly) in phase space. Hochman et al (2019) found that the dynamical systems persistence may be directly linked to the more conventional notion of persistence of weather regimes.

We consider a dynamical system with an observed trajectory $x(t)$, and a point ζ in phase space along this trajectory. Taking as example a daily SLP dataset, $x(t)$ would be a succession of daily SLP latitude-longitude maps, and ζ would be an SLP latitude-longitude map for a specific day (Fig. 1a). We are interested in the behavior of the system near ζ. To characterise this, we define a logarithmic return as:

$$g(x(t),\zeta) = -\log[\mathrm{dist}(x(t),\zeta)]. \tag{1}$$

Where dist is a distance function between two vectors, which tends to zero when x and y are close to each other. The values of $g(x(t),\zeta)$ define a time series of logarithmic returns which takes large values when the distance between $x(t)$ and ζ is small. Given $s(q,\zeta)$ a high q-th quantile of $g(x(t),\zeta)$, we introduce the exceedances $u(t,\zeta) = g(x(t),\zeta) - s(q,\zeta)$ with the condition $g(x(t),\zeta) > s(q,\zeta)$. The latter condition defines a recurrence, namely a case when the trajectory $x(t)$ approaches ζ (Fig. 1b). Going back to the SLP example, these would be days when the SLP map resembles closely that of the reference state ζ. The cumulative probability distribution of $F(u,\zeta)$ then converges to the exponential member of the Generalized Pareto Distribution (Freitas et al, 2010; Lucarini et al, 2012):

$$F(u,\zeta) \simeq \exp\left[-\vartheta(\zeta)\frac{u(\zeta)}{\sigma(\zeta)}\right]. \tag{2}$$

The parameters ϑ and σ depend on the point ζ chosen on the attractor. This means that, for our series of daily SLP latitude-longitude maps, each day will have a different value of ϑ and σ. As discussed below, ϑ and σ are fundamental quantities to describe the dynamical properties of the system. When applying this approach to a given dataset, all the maps in the data series serve in turn as reference point ζ.

2.2 Local dimension

The local dimension $d(\zeta)$ is obtained via the simple relation $d(\zeta) = 1/\sigma(\zeta)$. A short MatLab code to compute d is provided in Appendix B of Faranda et al (2019b). This result holds when $x(t)$ contains the system's full set of phase-space variables. In this case, d is independent of the chosen dist for all ζ. If $x(t)$ is an *observable* of the system, i.e. some smooth function of a variable

in the full phase space, $d(\zeta)$ can still be computed in this way (Rousseau and Saussol, 2010; Rousseau, 2014), but the quantitative results depend on the distance definition.

For practical reasons, here we use the Euclidean distance, which is also used to compute circulation analogues (e.g. Yiou et al (2013)). We also emphasize that the value of d is bounded by the number of coordinates of the observable and the dynamics of the system.

Given two observables (or two sets of observables) $x(t)$ and $y(t)$ of a "larger" system, we can define $d_x(\zeta)$ and $d_y(\zeta)$ (from now on we will drop the dependence on ζ). They are the dimensions of the Poincaré sections defined by x and y around ζ, with respect to the chosen dist. This would be equivalent to, for example, computing daily values of d for SLP and T2M fields separately, over a given geographical domain. We can further consider the Poincaré section jointly spanned by x and y. A state on this section is defined by the pair $\zeta = \{\zeta_x, \zeta_y\}$. The joint logarithmic returns can then be defined as:

$$g(x(t), y(t)) = -\log\left[\mathrm{dist}\left(\frac{x(t)}{||x||}, \frac{\zeta_x}{||x||}\right)^2 + \mathrm{dist}\left(\frac{y(t)}{||y||}, \frac{\zeta_y}{||y||}\right)^2\right]^{\frac{1}{2}} \quad (3)$$

Here, $||.||$ is the average root mean square norm of the coordinates of a vector. For example, $\|x\| = E_t\left(\left[\sum_i^K x_i(t)^2\right]^{\frac{1}{2}}\right)$, where K is the number of components of x and E_t is an average over time t. Based on Eq. (3), one can then compute $d_{x,y}$, namely the *co-dimension* between x and y. This is equivalent to considering the SLP and T2M fields jointly, and requiring that a recurrence – a day similar to the day of interest ζ – is such in both fields simultaneously.

The co-dimension can be used as a first quantity to characterise the mutual dependence of two observables. For two observables x and y of the system, the following properties hold:

$$\min(d_x, d_y) \leq d_{x,y} \leq d_x + d_y \quad (4)$$

where the min domain consists of the dimensions computed $\forall\ t$. The fact that the co-dimension is bounded by the sum of the individual dimensions suggests that it is a measure of the coupling of the two observables, namely the dimension required to describe their joint local behaviour. If x and y are uncoupled, then $d_{x,y} = d_x + d_y$. If x and y are linked with a deterministic function, then $d_{x,y} = \min(d_x, d_y)$. We exemplify these behaviours for a discrete dynamical system in Sec. 3.2.

2.3 Local persistence

The persistence of the state ζ is measured via the extremal index $0 < \vartheta(\zeta) < 1$, an adimensional parameter, from which we extract $\theta(\zeta) = \vartheta(\zeta)/\Delta t$. $\theta(\zeta)$ is

therefore the inverse of the average residence time of trajectories around ζ in units of frequency (in this study $\Delta t = 1$ day). If ζ is a fixed point of the attractor then $\theta(\zeta) = 0$ (trajectories stay at ζ). For a point that immediately leaves the neighborhood of ζ, then $\theta = 1$. Intermediate values of θ are obtained for regions of the attractor that are close to fixed points, and that trajectories leave "slowly". To estimate θ, we adopt the Süveges estimator (Süveges, 2007). Returning to our previous example, we consider all days on which the SLP map resembles the reference SLP state ζ, and we estimate the average number of consecutive days such a resemblance lasts for. Since we are investigating meteorological time scales, the choice of daily fields is appropriate to compute θ. Different applications might necessitate the use of a smaller (e.g. in the study of convective systems) or a larger (e.g. in the study of ocean dynamics) Δt.

As for d above, the procedure can be extended to more than one variable, such that one may define the inverse co-persistence $\theta_{x,y}$. The physical interpretation for atmospheric fields is analogous to that of the inverse persistence above, except that we now consider two fields jointly. Taking again SLP and T2M fields as examples, we thus consider series of consecutive days on which both resemble the respective reference states ζ_{SLP} and ζ_{T2M}. $\theta_{x,y}$ is a weighted average of θ_x and θ_y, where the weights depend on the density of observed neighbours of ζ. For a more detailed discussion we refer the reader to Abadi et al (2018).

2.4 Local co-recurrence ratio

Given two observables x and y, we define the co-recurrence ratio $0 \leq \alpha(\zeta) \leq 1$ of a state $\zeta = \{\zeta_x, \zeta_y\}$ as:

$$\alpha(\zeta) = \frac{\Pr\left[g(x(t)) > s_x(q) | g(y(t)) > s_y(q))\right]}{\Pr\left[g(x(t)) > s_x(q)\right]} \tag{5}$$

In terms of the SLP and T2M latitude-longitude fields, we compute the probability of having a day when the SLP map resembles that of the reference state ζ_{SLP}, given that the T2M map resembles that of the reference state ζ_{T2M}, versus the probability of only the SLP map resembling the relevant reference state. The thresholds $s_x(q)$ and $s_y(q)$ determine how close the monovariate recurrences have to be. Whenever x and y do not have the same units, a normalization $x/||x||$ and $y/||y||$ must be performed before computing α. When $\alpha(\zeta) = 0$, there are no co-recurrences of $\zeta = \{\zeta_x, \zeta_y\}$ when we observe a recurrence of ζ_x. When $\alpha(\zeta) = 1$, all the co-recurrences of $\zeta = \{\zeta_x, \zeta_y\}$ also correspond to recurrences of ζ_x.

From Bayes' theorem on conditional probabilities, we note that α does not depend on the order of x and y, which can hence be exchanged. This implies that α cannot be interpreted in terms of causation. We further note that α is not necessarily correlated with the co-dimension $d_{x,y}$, as discussed further in Sect 3.2. Finally, we underscore that α cannot be directly compared to

statistical dependence measures because of its local nature in phase space. Here, we have defined α in the case of two variables x and y, but the approach can be extended to more complex multivariate cases.

3 Application to stochastic processes and dynamical systems

3.1 Gaussian Bivariate copula

To illustrate the interpretation of α, we draw data from random Gaussian bivariate copula distributions, namely:

$$\phi(z, \mu, \Sigma) = \frac{1}{\sqrt{4\pi^2|\Sigma|}} \exp\left[-\frac{1}{2}(z-\mu)'\Sigma^{-1}(z-\mu)\right] \tag{6}$$

where $z = (x, y)$ is an independent and identically distributed (i.i.d.) variable vector in $[0, 1] \times [0, 1]$, μ is the mean vector, Σ is a 2×2 symmetric matrix and $|\Sigma|$ is its norm. The diagonal elements of Σ contain the variances for each variable (x, y), while the off-diagonal elements of Σ contain the covariances between variables (x, y). In the example we set $\mu = \{0, 0\}$ and the diagonal elements of Σ equal to one. We term the non–diagonal elements ρ.

By varying ρ, we obtain a range of different behaviours. There are two limiting cases: for $\rho = 0$, $\{x, y\}$ the data cover a large portion of the unit square; for $\rho = 1$, $x = y$ and the data are aligned along the first diagonal. We illustrate graphically four different cases: $\rho = 0.5$, $\rho = 0.75$, $\rho = 0.9992$, and $\rho = 1$. We generate 10^5 values of z for each case. For clarity we report 2×10^3 points in Figure 2. These are colored by the values of α, the co-recurrence ratio, computed by fixing the quantile $q = 0.98$.

α can be interpreted as follows: when the variables are locally independent (Figure 2a, b), $\alpha = 0$ almost everywhere. When ρ is close to 1 (Figure 2c), near the edges $(0, 0)$ and $(1, 1)$ the dependence is very strong and α=1. Near the centre $(x = y = 0.5)$ the dependence is lower and $\alpha < 1$. This behavior is determined by the very nature of the distribution considered here, which is constrained to be close to the diagonal near the edge points. In the perfect coupling case $\rho = 1$ (Figure 2d), $x = y$ and $\alpha = 1$ along the full length of the diagonal.

3.2 Baker's map

To further illustrate the properties of the dynamical indicators, we analyse a modified version of baker's map, defined on the unit square $[0, 1] \times [0, 1]$ iteratively for $n \geq 1$, by:

$$x_{n+1} = f(x_n) = \begin{cases} ax_n, & \text{for } y_n < c \\ \frac{1}{2} + bx_n, & \text{for } y_n > c \end{cases} \tag{7}$$

and

$$y_{n+1} = f(y_n) = \begin{cases} y_n/a, & \text{for } y_n < c \\ (y_n - c)/(1-c), & \text{for } y_n > c \end{cases} \quad (8)$$

with $c = 1/3$, $a = 1/5$ and $b = 1/4$. For these parameter values, the map is expanding on x and contracting on y. Figure 3a shows the first 10^5 iterations of the map. Along the x direction, the map consists of a set with dimension $d_x < 1$ while along y there is a collection of lines yielding a dimension $d_y = 1$.

We can add coupling between x and y as:

$$x_{n+1} = (1-\epsilon)f(x_n) + \epsilon f(y_n)$$

with $0 \leq \epsilon \leq 1$. When ϵ is zero, x and y are uncoupled. As the noise intensity ϵ increases, the dynamics of x, y becomes more and more synchronized. Finally, when $\epsilon = 1$, the dynamics is synchronized and concentrated on the diagonals (Figure 3b).

To illustrate the properties of the co-dimension and the co-recurrence coefficients in this set-up, we change the values of the noise intensity ϵ in the range $10^{-10} < \epsilon < 1$ and perform 10^6 iterations of the maps starting from random initial conditions on the unit square. The first 10^4 iterations are discarded. We then compute the dynamical indicators at 500 points ζ. Figure 4 shows the results of this computation for the dimensions d_x, d_y, the co-dimension $d_{x,y}$ and the co-recurrence coefficient α. In the limit $\epsilon \to 0$, $d_x + d_y = d_{x,y}$ and $\alpha \simeq 0$. When increasing the coupling, $d_{x,y} < d_x + d_y$. In the limit for large noise intensity $\epsilon = 1$, $x \propto y$, $d_{x,y} = 1$ and the co-recurrence coefficient α approaches 1. The relatively sharp transition happening for noise intensity $\epsilon \simeq 10^{-3}$ can be explained by drawing from the behavior of stochastically perturbed systems (Faranda et al, 2013). In the latter, when noise spanning several orders of magnitude is added, it has an appreciable effect only when stochastic fluctuations (of order ϵ) have the same order as those induced by the deterministic dynamics. If the noise fluctuations are larger than those of the deterministic system, the dimension d reflects the stochastic system.

This simple example highlights the difference in the information provided by the co-dimension and co-recurrence ratio. The co-dimension generally decreases for increasing coupling, while the co-recurrence ratio increases. We further note that the co-recurrence ratio provides a normalised coupling measure, while the value of the co-dimension depends on the relative dimension on the x and y manifold. The behavior for θ is not shown as, for this system, we have trivially that $\theta_x = \theta_y = \theta_{x,y}\ \forall\ \epsilon$.

4 Application to North American temperature extremes

4.1 Data and Methods

We base our study on NCEP/NCAR reanalysis data (Kalnay et al, 1996) over the period 1948-2018, with a horizontal resolution of 2.5°. We consider a domain spanning North America (170°W≤Long.≤40°W, 22.5°N≤Lat.≤70°N).

We adopt daily SLP as the meteorological variable to describe the large scale atmospheric circulation. A wealth of atmospheric features, ranging from teleconnection patterns to storm track activity to atmospheric blocking can be diagnosed from the SLP field (e.g. Murray and Simmonds (1991); Yiou et al (2013); Comas-Bru and McDermott (2014)). We diagnose temperature extremes using T2M. Throughout the analysis, we consider winter (December-January-February: DJF) and summer (June-July-August: JJA) seasons separately.

We define warm or cold extremes following the procedure of Messori et al (2018). This amounts to using area-averaged temperature anomalies relative to a running daily climatology and applying some smoothing. Here, we consider the domain 100°W≤Long.≤70°W, and 30°N≤Lat.≤45°N. This corresponds to a densely populated part of western North America (see Figure 5). We then define two anomaly distributions for each season: one for all days with negative area-averaged temperature anomalies and one for all days with positive area-averaged temperature anomalies. The percentiles discussed in the text below are relative to these distributions. We specifically define cold days as days below the 10th percentile of the negative-anomalies distributions and warm days as days above the 90th percentile of the positive-anomalies distributions. We refer to these throughout the text as *warm days* or *warm extremes* and *cold days* or *cold extremes*, respectively. To avoid double-counting a single warm or cold event, following Messori et al (2016) we impose a minimum separation between successive warm or cold days – here set at one week. All geographical plots show single-gridpoint T2M anomalies defined relative to a daily seasonal cycle, obtained by averaging all days available in the dataset. For example, the climatological value at a given gridbox for the 3rd December is the average value at that gridbox for all 68 3rd Decembers in the data.

4.2 Dynamical properties and seasonality

We analyse the relations between local dimension d, persistence θ^{-1} and co-recurrence ratio α for SLP and T2M. We consider the cases where d and θ are computed on each variable individually and on both variables at the same time. α is naturally always computed on both variables. The results are reported in Figure 6 in the form of dimension-persistence diagrams. Each point in the diagram represents a pair (d, θ) corresponding to the patterns observed on a given day for the variable(s) of interest. The colours show the values of α.

α is negatively correlated with d_{SLP} (Figure 6a), meaning that configurations with lower d typically favour a higher coupling. The fact that d and θ are themselves correlated implies that the latter configurations are, on average, also highly persistent. The picture for T2M is less clear (Figure 6b): strongly coupled states often correspond to high d and low-persistence configurations. Finally, the joint analysis highlights a similar pattern as for SLP, with α showing a clear dependence on the co-dimension $d_{SLP,T2M}$.

Faranda et al (2017a) and Rodrigues et al (2018) found that the correlations evident in dimension-persistence diagrams, as those shown in Figure 6, reflected the strong seasonal dependence of the metrics. We therefore investigate the seasonal cycle in the above quantities (Figure 7). All three indicators are subject to a marked seasonality, with α presenting a clear peak in summertime, when mean values are almost twice those of the rest of the year. The local dimension d instead displays maxima in the shoulder seasons for SLP, and in summer for T2M. The broad summertime maximum in d_{T2M} closely matches the timing of the α peak; d_{SLP} displays a minimum at roughly the same time. This explains the very different relations between d and α shown in Figures 6a and b for SLP and T2M, respectively. The variability in local dimensions is not, however, as marked as that seen in α. The co-dimension $d_{SLP,T2M}$ does not follow the seasonal cycle of either d_{SLP} or d_{T2M}. It peaks in late spring/early summer and early autumn, while displaying mid-summer and winter minima. θ_{SLP} is mostly in phase with d_{SLP}, as is θ_{T2M} with d_{T2M}, albeit with a larger seasonal variability. The inverse co-persistence $\theta_{SLP,T2M}$ again does not follow either of the single-variable cycles, but rather displays an oscillatory behavior throughout the year, somewhat reminiscent of the variability in $d_{SLP,T2M}$. The largest values tend to occur during the autumn and winter months, while the lowest values are mostly found in spring.

The seasons during which most of the highest and lowest values of the three indicators occur reflect these seasonal cycles (Figure 8). Even though the mean of α has a very clear summertime peak, the variability peaks instead in winter (Figure 7a), implying that a small number of α maxima are seen during this season. This matches the timing of minima in both d_{SLP} and θ_{SLP}, consistent with what shown in Figure 6a. The α minima occur instead predominantly in winter and spring, when d_{SLP} and θ_{SLP} display relatively large values. Analogous conclusions can be drawn concerning the seasonal cycles of the co-dimension and co-persistence (Figure 8).

4.3 Temperature extremes

The above results show that d, θ and α display both a large intra-seasonal variability and a strong seasonal cycle. The most striking features are the distinct wintertime and summertime behaviours in α. We ascribe the seasonality of the three metrics to the fact that they reflect the very different atmospheric dynamics found in the different seasons (see Faranda et al (2017b,a)). Even though the mean values of the metrics and an analysis of their seasonal variability can provide interesting physical insights, such as assessing which seasons present more or less coupling between SLP and T2M, the real added value in having instantaneous metrics is the possibility of using them to study climate extremes. Here, we specifically investigate warm and cold extremes over North America, during both the summer (JJA) and winter (DJF) seasons.

Summer extremes. We begin by considering summertime warm and cold days (see Section 4.1). The co-dimension and co-persistence, as a function of the temperature percentile over the selection domain, show little change for extreme warm and extreme cold events relative to their respective seasonal cycles (continuous lines in Figure 9a,b). Similarly, α (Figure 9c) remains close to its seasonal cycle, regardless of the temperature percentile. This points to the fact that the joint d and θ metrics do not reflect the evolution of the temperature field and that, consistently with this, the co-recurrence ratio is also largely insensitive to temperature extremes. The d and θ metrics computed for the SLP and T2M fields individually (dashed and dashed-dotted lines in Figure 9a,b) display a similar behaviour.

To highlight the link between the dynamical systems metrics and the underlying physical properties of the atmospheric circulation associated with different events, we define an average daily baroclinc vector as:

$$B = \overline{\nabla(SLP) \times \nabla(T2M)} \tag{9}$$

where $\times$ denotes the vector product and the overbar a spatial mean over the domain highlighted in black in Figure 5. B is in units of [hPa$\times$K]/[$^\circ$lon$\times^\circ$lat]. Values close to zero indicate a predominantly barotropic atmosphere i.e. the isolines of pressure are aligned with those of temperature. Positive and negative values correspond instead to warm and cold air advection, respectively. Two caveats are that we do not account for cancellation between opposing advections in the domain, and may give excessive weight to zonal advection at northerly latitudes. Domain-averaged values for B are shown in Figure 9d. In summer, all the values of B are generally close to zero. This decoupling of the temperature anomalies from specific large-scale advection patterns is also reflected in the T2M and SLP composites for α extremes. Regardless of whether coupling maxima or minima are chosen, the anomalies are weak throughout the domain, and the sign agreement between the different extreme α occurrences is low (Figure 10).

Winter extremes. The situation in winter is radically different: $d_{SLP,T2M}$, $\theta_{SLP,T2M}$ and α for the cold extremes all show significant deviations from the climatology. The warm extremes show weak deviations in $d_{SLP,T2M}$ and $\theta_{SLP,T2M}$, and a moderate deviation in α. The signs of the d and θ deviations for the two sets of temperature extremes are opposite: cold extremes show negative anomalies, while the warm extremes show weak positive anomalies (continuous lines in Figure 11a-c). The d and θ computed for the two variables individually show a similar behaviour to the bivariate indicators for the cold extremes (blue dashed and dashed-dotted lines in Figure 11a, b). For the warm extremes, d_{SLP} and θ_{SLP} are broadly consistent with $d_{SLP,T2M}$ and $\theta_{SLP,T2M}$ (red dashed lines in Figure 11a, b), while d_{T2M} and θ_{T2M} show anomalies of a different sign to those of the bivariate indicators (red dash-dotted lines in Figure 11a, b). However, they still show a clear deviation from the climatology. The co-recurrence ratio also shows deviations of the opposite

sign for the two sets of extremes: negative for warm extremes and positive for cold extremes (Figure 11c). The analysis of the baroclinic vector B (Figure 11d) provides clear insights on such asymmetry: cold anomalies systematically correspond to a strong cold advection ($B < 0$), reflecting the high value of the co-recurrence ratio α, while the warm extremes display a significant warm advection ($B > 0$) only for the highest percentiles.

By conditioning purely on low and high values of α, one recovers strong temperature anomalies over the domain of interest (Figure 12b, d). This is particularly evident for the high α, which correspond to cold anomalies across a broad swath of the North American continent. Both low and high α are associated with anomalous SLP patterns favoring strong meridional advection (Figure 12a, c). The SLP and T2M fields therefore appear to display recurring joint large-scale configurations which favour cold extremes and, to a lesser extent, warm extremes. This may seem in contradiction with the information provided by α, as one may expect no coherent large scale SLP patterns to be associated with α minima (and hence warm extremes). We expand on this aspect in the discussion section below.

5 Discussion and conclusions

We have proposed two metrics to diagnose the properties of instantaneous configurations of atmospheric variables, namely local dimension d and inverse persistence θ. The first is a proxy for the number of active degrees of freedom of a given atmospheric configuration, while the second measures the average time over which a configuration is maintained. These indicators have been previously applied to individual atmospheric variables (e.g. Faranda et al (2016, 2017b,a, 2019c); Messori et al (2017); Rodrigues et al (2018); Scher and Messori (2018, 2019); Hochman et al (2019)). Here, we present their use in a multivariate context for both simple dynamical systems and climate reanalysis data. We further introduce a novel metric which measures the co-occurrence of analogue states of two (or more) variables: the co-recurrence ratio α. A small α indicates rare co-recurrences, while a large α indicates that recurrences of a given state in one of the two variables systematically match recurrences of a corresponding state in the other. This metric thus elucidates the mutual dependence between the dynamical evolutions of the two variables. Unlike statistical dependence measures, α is grounded in the topology of the variables' phase-space, and therefore in the local geometry of the underlying attractor. Locality in phase space, which translates to simultaneity in time, makes the above framework well-suited to the study of atmospheric variability and extremes. From the numerical analysis performed in this work on simple dynamical systems and climate data, we infer that the co-recurrence ratio is an indicator of physical coupling and underscores interactions between variables. This understanding comes from the observation that α relates non-trivially to the co-dimension (cf. the baker's map and atmospheric data). The results obtained for atmospheric data suggest that, when α is high, the phase space

around ζ_x and ζ_y shrinks yielding lower co-dimensions. In other words, when α is high the variables are dynamically dependent. We expect this behaviour to be representative of high-dimensional systems.

As an illustration on atmospheric data, we have investigated the occurrence of cold or warm days in North America, linking temperature anomalies over the eastern part of the continent to the larger-scale T2M and SLP fields. We identify strongly and weakly coupled states in these two variables, and then elucidate the interplay between these and the occurrence of the temperature extremes.

Summertime temperature extremes appear to be unrelated to the joint dynamical properties of the T2M and SLP large-scale fields. This is also reflected in the d and θ metrics computed on the individual variables and in the weak co-recurrence between SLP and T2M. The recurrences of T2M fields associated with temperature extremes are not matched by recurrences of the same SLP fields, or at least not more so than what is typical for any day in the summer season. Indeed, the sign agreement of the large-scale SLP anomalies associated with the warm and cold temperature extremes are relatively weak and display very limited sign agreement (see Fig. 13 for the upper/lower 10 percentiles). Similarly, a proxy for temperature advection shows near-zero values (B in Fig. 9d. This points to local factors, not captured by the large-scale SLP field analysed here, as contributors to these events (e.g. local soil moisture anomalies, small-scale precipitation events, local sensible heat fluxes etc.).

In winter, the picture is different (see Fig. 14 for the SLP anomalies corresponding to the upper/lower 10 percentiles of T2M anomalies). Both the joint and monovariate d and θ discriminate warm and cold extremes, and these also display anomalous α values. Cold extremes are characterised by an anomalously high persistence, low local dimension and strong coupling. This points to the need for persistent circulation patterns and large-scale cold advection (Fig 14c) for a cold day to occur, as well as to the fact that whenever cold days occur, similar large-scale T2M and SLP patterns are found (Cellitti et al, 2006; Grotjahn, 2016; Walsh et al, 2001; Messori et al, 2016). Here we identify these as: a large anticyclone over North America, advecting cold, dry air from the Arctic region (as reflected by the baroclinic vector B in Fig. 11d), and two low-pressure cores on either flank (Fig. 14c). Such pattern leads to widespread low temperatures as far south as Texas and Northern Florida (Fig. 14d). Warm extremes are instead characterised by weak negative anomalies in persistence, weak positive anomalies in local dimension and a below-climatology coupling. The large-scale pattern is roughly inverse to that seen for the cold days, with a cyclonic anomaly over the continent favouring advection of warm, moist air from the low latitudes (Figs. 11d and 14a) and leading to high temperatures on the eastern part of the North American continent (Fig. 14b). Wintertime warm days are thus associated with a clear large-scale SLP pattern, yet the SLP and T2M fields do not necessarily display frequent joint recurrences, this might mean that warm winter days are not dynamically different from the other "average" days of the season, contrary to cold episodes.

There is an evident ambiguity in an apparently coherent large scale anomalous SLP pattern and the claim of a weak coupling to the T2M field. A possible explanation is that this issues from the dynamical systems analysis being performed on the absolute fields rather than on the anomalies shown in the figures. Since the co-recurrence ratio is sensitive to the seasonal cycle (Figure 7a), this may introduce a discrepancy between absolute and anomalous recurrence patterns. This points to the need to perform a systematic analysis on the anomaly fields, although this is far from trivial in a dynamical systems context. There is no currently accepted framework for removing the average attractor components and then characterising the residual fluctuations. Formally, this would correspond to analysing recurrences within hyper-ellipsoids whose the eccentricity depends on the position in phase space, rather than hyper-spheres.

Our analysis has shown insights that can be provided by our local (in phase-space, instantaneous in time) and multi-variate dynamical systems analysis as well as some ongoing challenges. As such, it should be viewed as a proof of concept for the utility of this approach. More generally, there is no requirement for the different variables to be geographically co-located or temporally coincident, as was the case here. For example, one could select SLP over a region upstream of the target temperature region, or could use SLP fields lagged by a few days relative to the T2M data. This would enable to study forcing instead of coupling between variables. The approach may also be used to study changes in the coupling and the interaction of variables under anthropogenic forcing, in a similar fashion as Faranda et al (2019a).

In this study we have considered bivariate applications and climate data on latitude-longitude grids. However the dynamical systems indicators can be generalised to higher dimensional variables or to a larger number of variables. We note, however, that in higher dimensional space we lose the x, y symmetry in the definition of the co-recurrence ratio, which then needs to be defined with respect to only one of the variables of the system. Our approach is thus flexible, and is applicable to observables of any complex system.

Acknowledgements DF, PY and GM acknowledge the support of the ERC grant No. 338965-A2C2. GM was partly supported by the Swedish Research Council Vetenskapsrådet No. 2016-03724. DF and PY were further supported by a CNRS-INSU LEFE/MANU grant (DINCLIC project).

References

Abadi M, Freitas ACM, Freitas JM (2018) Dynamical counterexamples regarding the extremal index and the mean of the limiting cluster size distribution. arXiv preprint arXiv:180802970

Barros V, Liao L, Rousseau J (2019) On the shortest distance between orbits and the longest common substring problem. Advances in Mathematics 344:311–339

Bevacqua E, Maraun D, Hobæk Haff I, Widmann M, Vrac M (2017) Multivariate statistical modelling of compound events via pair-copula constructions: analysis of floods in ravenna (italy). Hydrology and Earth System Sciences 21(6):2701–2723

Cellitti MP, Walsh JE, Rauber RM, Portis DH (2006) Extreme cold air outbreaks over the united states, the polar vortex, and the large-scale circulation. Journal of Geophysical Research: Atmospheres 111(D2)

Comas-Bru L, McDermott F (2014) Impacts of the ea and sca patterns on the european twentieth century nao–winter climate relationship. Quarterly Journal of the Royal Meteorological Society 140(679):354–363

Coumou D, Petoukhov V, Rahmstorf S, Petri S, Schellnhuber HJ (2014) Quasi-resonant circulation regimes and hemispheric synchronization of extreme weather in boreal summer. Proceedings of the National Academy of Sciences 111(34):12,331–12,336

Davies HC (2015) Weather chains during the 2013/2014 winter and their significance for seasonal prediction. Nature Geoscience 8(11):833

Eckmann JP, Ruelle D (1992) Fundamental limitations for estimating dimensions and lyapunov exponents in dynamical systems. Physica D: Nonlinear Phenomena 56(2-3):185–187

Faranda D, Freitas JM, Lucarini V, Turchetti G, Vaienti S (2013) Extreme value statistics for dynamical systems with noise. Nonlinearity 26(9):2597

Faranda D, Masato G, Moloney N, Sato Y, Daviaud F, Dubrulle B, Yiou P (2016) The switching between zonal and blocked mid-latitude atmospheric circulation: a dynamical system perspective. Climate Dynamics 47(5-6):1587–1599

Faranda D, Messori G, Alvarez-Castro MC, Yiou P (2017a) Dynamical properties and extremes of northern hemisphere climate fields over the past 60 years. Nonlinear Processes in Geophysics 24(4):713

Faranda D, Messori G, Yiou P (2017b) Dynamical proxies of north atlantic predictability and extremes. Scientific reports 7:41,278

Faranda D, Sato Y, Saint-Michel B, Wiertel C, Padilla V, Dubrulle B, Daviaud F (2017c) Stochastic chaos in a turbulent swirling flow. Physical review letters 119(1):014,502

Faranda D, Alvarez-Castro MC, Messori G, Rodrigues D, Yiou P (2019a) The hammam effect or how a warm ocean enhances large scale atmospheric predictability. Nature communications 10(1):1316

Faranda D, Messori G, Vannitsem S (2019b) Attractor dimension of time-averaged climate observables: insights from a low-order ocean-atmosphere model. Tellus A: Dynamic Meteorology and Oceanography 71(1):1–11

Faranda D, Sato Y, Messori G, Moloney NR, Yiou P (2019c) Minimal dynamical systems model of the northern hemisphere jet stream via embedding of climate data. Earth System Dynamics 10(3):555–567

Freitas ACM, Freitas JM, Todd M (2010) Hitting time statistics and extreme value theory. Probab Theor Rel 147(3-4):675–710

Galambos J, Lechner J, Simiu E, Hagwood C (1994) Extreme value theory for applications:. NIST technical publications

Gálfi VM, Bódai T, Lucarini V (2017) Convergence of extreme value statistics in a two-layer quasi-geostrophic atmospheric model. Complexity 2017

Gnedenko B (1943) Sur la distribution limite du terme maximum d'une serie aleatoire. Annals of mathematics pp 423–453

Grassberger P (1983) Generalized dimensions of strange attractors. Physics Letters A 97(6):227–230

Grassberger P, Procaccia I (1984) Dimensions and entropies of strange attractors from a fluctuating dynamics approach. Physica D: Nonlinear Phenomena 13(1-2):34–54

Grotjahn R (2016) Western north american extreme heat, associated large scale synoptic-dynamics, and performance by a climate model. Dynamics and Predictability of Large-Scale, High-Impact Weather and Climate Events 2:198

Harnik N, Messori G, Caballero R, Feldstein SB (2016) The circumglobal north american wave pattern and its relation to cold events in eastern north america. Geophysical Research Letters 43(20)

Hirschi M, Seneviratne SI, Alexandrov V, Boberg F, Boroneant C, Christensen OB, Formayer H, Orlowsky B, Stepanek P (2011) Observational evidence for soil-moisture impact on hot extremes in southeastern europe. Nature Geoscience 4(1):17

Hochman A, Alpert P, Harpaz T, Saaroni H, Messori G (2019) A new dynamical systems perspective on atmospheric predictability: Eastern mediterranean weather regimes as a case study. Science advances 5(6):eaau0936

Huke J (2006) Embedding nonlinear dynamical systems: A guide to takens' theorem

Kalnay E, Kanamitsu M, Kistler R, Collins W, Deaven D, Gandin L, Iredell M, Saha S, White G, Woollen J, et al (1996) The ncep/ncar 40-year reanalysis project. B Am Meteorol Soc 77(3):437–471

Kornhuber K, Osprey S, Coumou D, Petri S, Petoukhov V, Rahmstorf S, Gray L (2019) Extreme weather events in early summer 2018 connected by a recurrent hemispheric wave-7 pattern. Environmental Research Letters 14(5):054,002

Lucarini V, Faranda D, Wouters J (2012) Universal behaviour of extreme value statistics for selected observables of dynamical systems. Journal of statistical physics 147(1):63–73

Marwan N, Romano MC, Thiel M, Kurths J (2007) Recurrence plots for the analysis of complex systems. Physics reports 438(5-6):237–329

Messori G, Caballero R, Gaetani M (2016) On cold spells in north america and storminess in western europe. Geophysical Research Letters 43(12):6620–6628

Messori G, Caballero R, Faranda D (2017) A dynamical systems approach to studying midlatitude weather extremes. Geophysical Research Letters 44(7):3346–3354

Messori G, Woods C, Caballero R (2018) On the drivers of wintertime temperature extremes in the high arctic. Journal of Climate 31(4):1597–1618

644 Messori G, Davini P, Alvarez-Castro MC, Pausata FS, Yiou P, Caballero R
645 (2019) On the low-frequency variability of wintertime euro-atlantic plane-
646 tary wave-breaking. Climate Dynamics 52(3-4):2431–2450
647 Murray RJ, Simmonds I (1991) A numerical scheme for tracking cyclone cen-
648 tres from digital data. Australian Meteorological Magazine 39(3):155–166
649 Palmer T (2013) Climate extremes and the role of dynamics. Proceedings of
650 the National Academy of Sciences 110(14):5281–5282
651 Pickands III J (1975) Statistical inference using extreme order statistics. the
652 Annals of Statistics pp 119–131
653 Rodrigues D, Alvarez-Castro MC, Messori G, Yiou P, Robin Y, Faranda D
654 (2018) Dynamical properties of the north atlantic atmospheric circulation
655 in the past 150 years in cmip5 models and the 20crv2c reanalysis. Journal
656 of Climate (2018)
657 Rousseau J (2014) Hitting time statistics for observations of dynamical sys-
658 tems. Nonlinearity 27(9):2377
659 Rousseau J, Saussol B (2010) Poincaré recurrence for observations. Transac-
660 tions of the American Mathematical Society 362(11):5845–5859
661 Scher S, Messori G (2018) Predicting weather forecast uncertainty with
662 machine learning. Quarterly Journal of the Royal Meteorological Society
663 144(717):2830–2841
664 Scher S, Messori G (2019) Weather and climate forecasting with neural net-
665 works: using general circulation models (gcms) with different complexity as
666 a study ground. Geoscientific Model Development 12(7):2797–2809
667 Süveges M (2007) Likelihood estimation of the extremal index. Extremes 10(1-
668 2):41–55
669 Von Storch H, Zwiers FW (2001) Statistical analysis in climate research. Cam-
670 bridge university press
671 Walsh JE, Phillips AS, Portis DH, Chapman WL (2001) Extreme cold
672 outbreaks in the united states and europe, 1948–99. Journal of climate
673 14(12):2642–2658
674 Yiou P, Salameh T, Drobinski P, Menut L, Vautard R, Vrac M (2013) Ensem-
675 ble reconstruction of the atmospheric column from surface pressure using
676 analogues. Climate dynamics 41(5-6):1333–1344
677 Zscheischler J, Seneviratne SI (2017) Dependence of drivers affects risks asso-
678 ciated with compound events. Science Advances 3(6):e1700,263

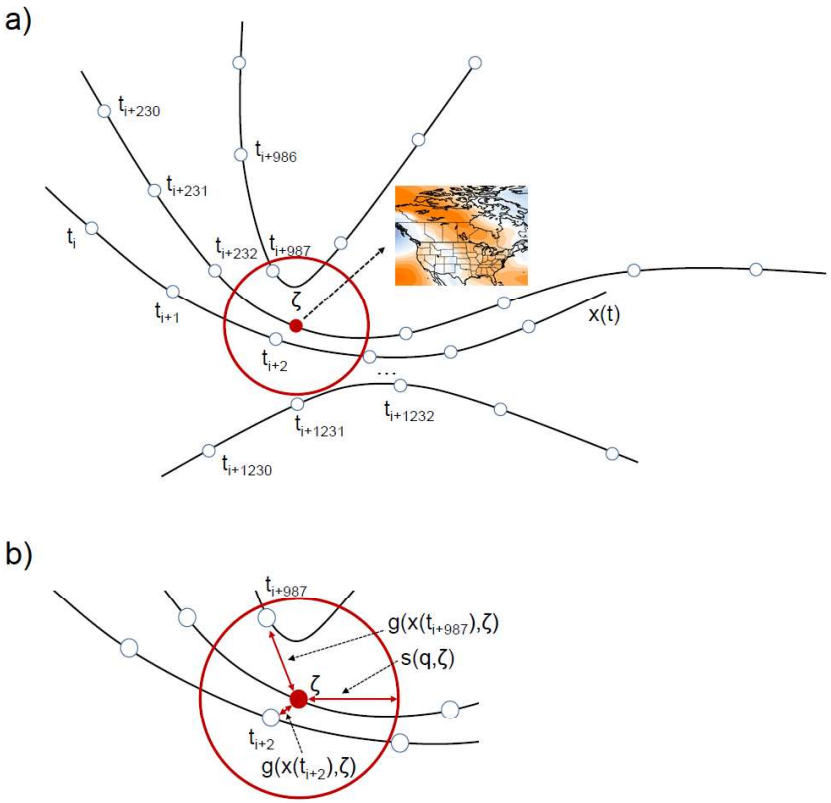

Fig. 1 (a) Idealised schematic of a long trajectory $x(t)$ representing the evolution of a sea-level pressure field. Note that all segments shown in the panel are part of a single, long trajectory. The white circles along the trajectory represent discrete, instantaneous measurements of the continuous evolution of the field, such as may be provided by a reanalysis dataset. The red point ζ is the state of interest, for example a specific sea-level pressure map at a specific time. (b) Visual illustration of the hyper-sphere determined by the high threshold $s(q, \zeta)$, which defines recurrences, and the distances between measurements defined by $g(x(t), \zeta)$. Note that g is defined so as to take large values for small separations. Thus $g(x(t), \zeta) > s(q, \zeta)$ holds for all points within the hyper-sphere. In the idealised schematic, only two measurements satisfy this condition.

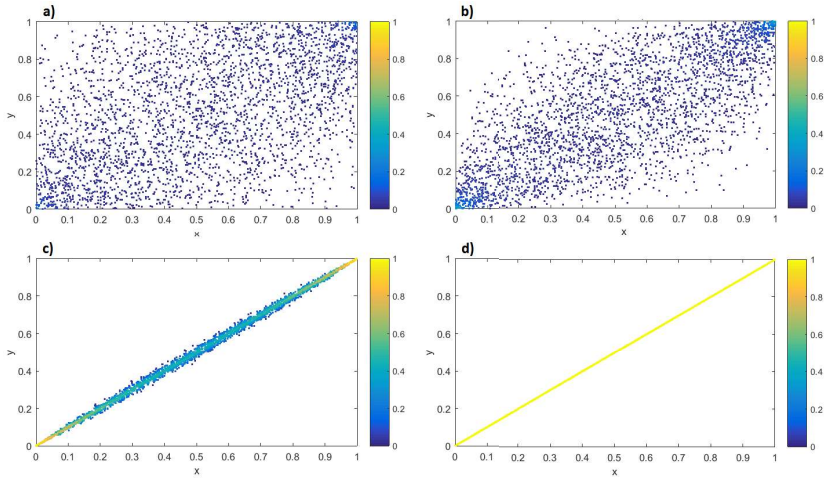

Fig. 2 Four different realizations of the Gaussian bivariate copula (Eq. 6) with $\rho = 0.5$ (a), $\rho = 0.75$ (b), $\rho = 0.9992$ (c), $\rho = 1$ (d). The co-recurrence ratio α is displayed in colours.

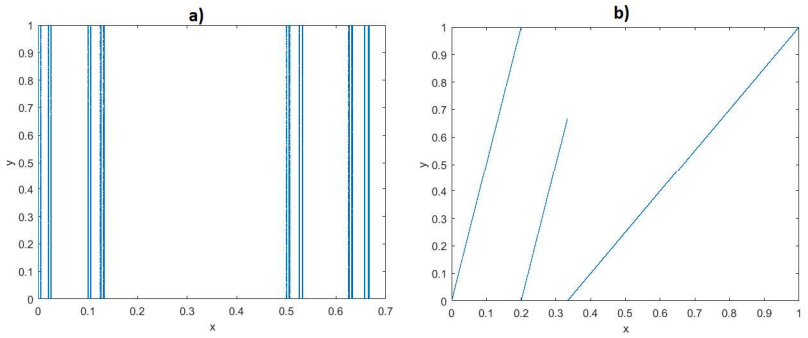

Fig. 3 Two different realizations of 10^5 iterations of the baker's map (Eqs. 7-8) with $\epsilon = 0$ (a), and $\epsilon = 1$ (b).

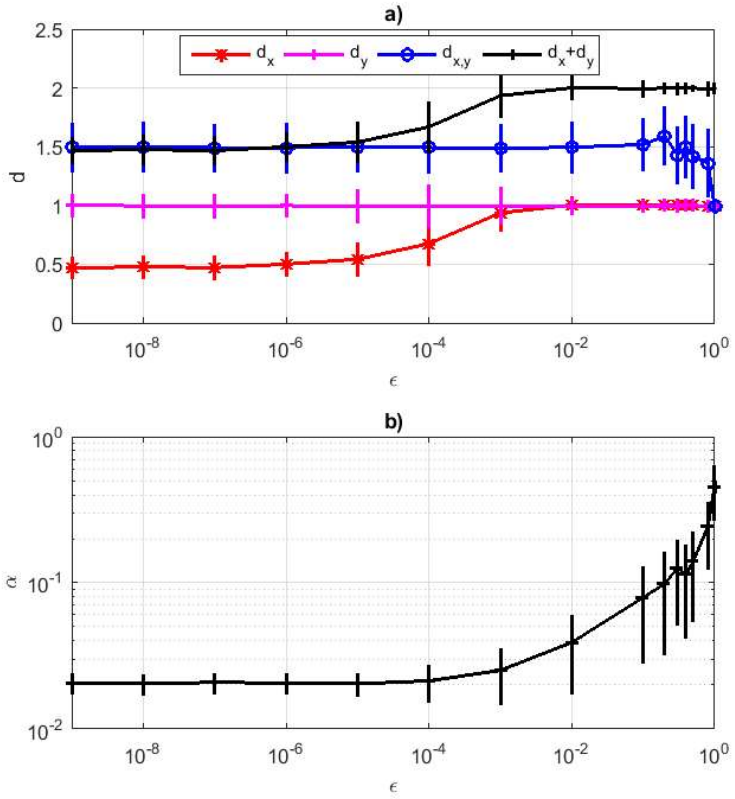

Fig. 4 Dynamical indicators for the baker's map (Eqs. 7-8). a) Dimensions d_x (red line, asterisks) and d_y (pink line, no markers), co-dimensions $d_{x,y}$ (blue line, asterisks) and d_x+d_y (black line, no markers) for different values of ϵ. b) Co-recurrence ratio α for different values of ϵ. Error bars show one standard deviation of the d, α distributions.

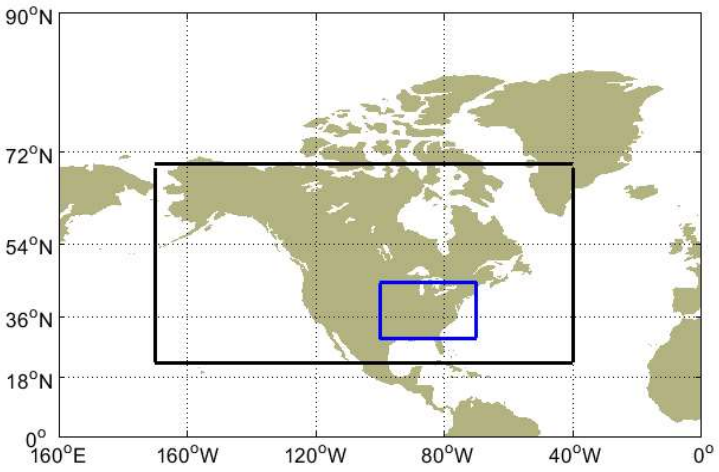

Fig. 5 The region studied in this article. The black rectangle highlights the domain over which the dynamical indicators are computed; the blue rectangle highlights the domain over which the temperature anomalies are computed. Note that only land points are taken into account in the area-averaged temperature anomalies over the domain.

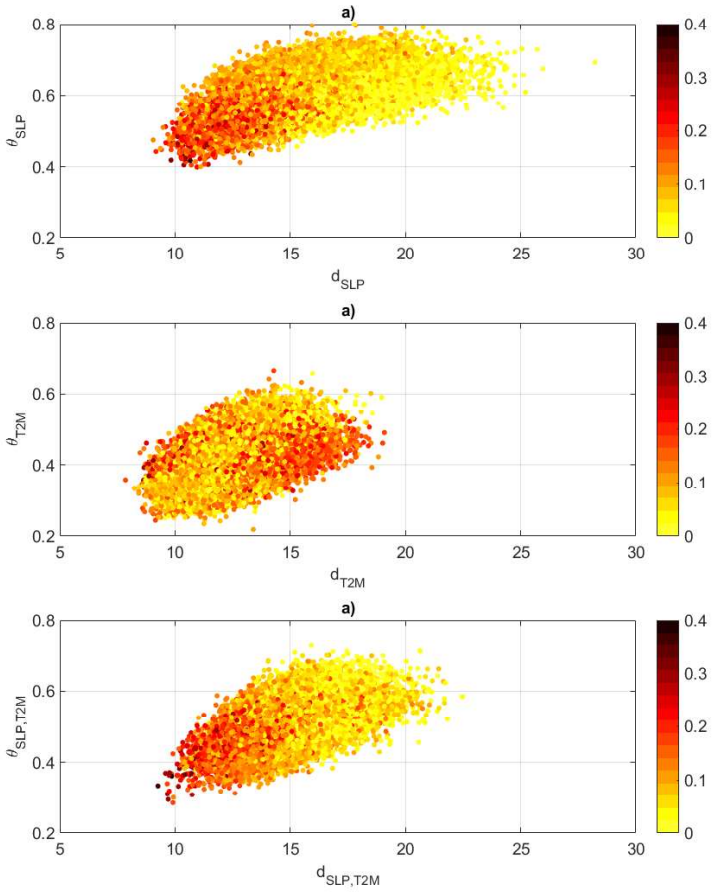

Fig. 6 Local dimension d and inverse persistence θ scatter plots coloured with the values of the co-recurrence ratio α. a) d_{SLP} and θ_{SLP}. b) d_{T2M} and θ_{T2M}. c) $d_{SLP,T2M}$ and $\theta_{SLP,T2M}$.

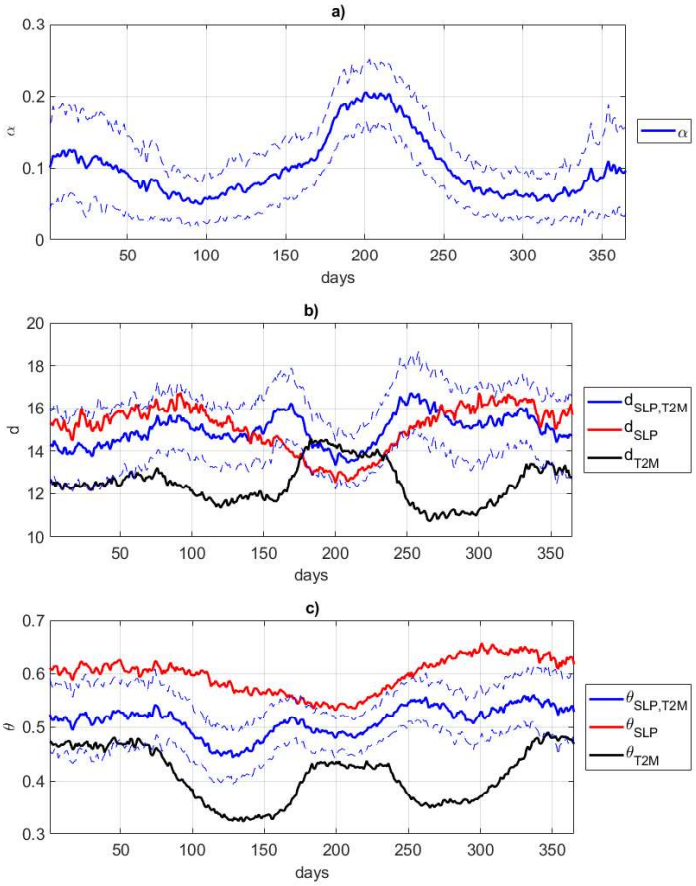

Fig. 7 Seasonal cycles of co-recurrence ratio α (a), local dimension d (b) and inverse persistence θ (c). The continuous red, black and blue lines in panels (b, c) correspond to d, θ computed on SLP, T2M and both variables jointly, respectively. The dashed lines mark one standard deviation of the quantities represented by the blue lines.

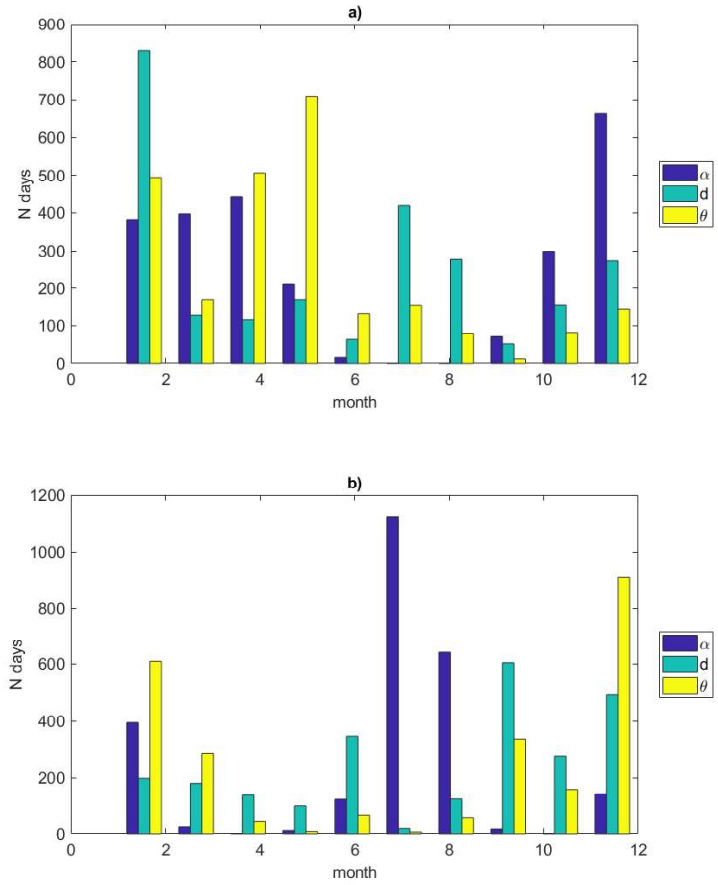

Fig. 8 Number of minima (a) and maxima (b) of co-recurrence ratio α, local co-dimension $d_{SLP,T2M}$ and inverse co-persistence $\theta_{SLP,T2M}$ per month. These are defined as the 10% highest and lowest values of the relevant distributions, respectively.

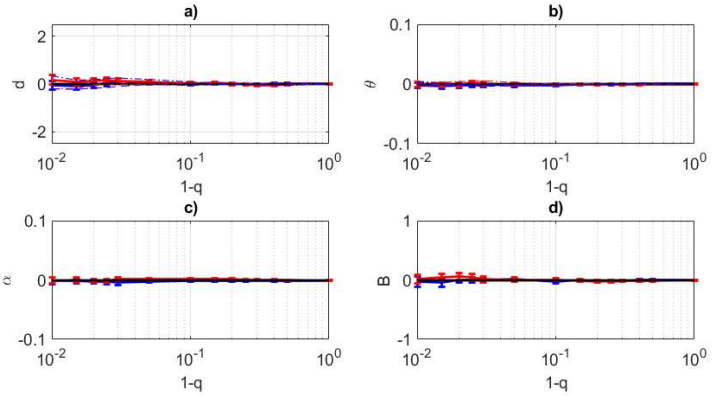

Fig. 9 Local dimensions d_{SLP} (dashed), d_{T2M} (dashed-dotted) and co-dimension $d_{SLP,T2M}$ (continuous) (a); inverse persistences θ_{SLP} (dashed), θ_{T2M} (dahsed-dotted) and inverse co-persistence $\theta_{SLP,T2M}$(continuous) (b); co-recurrence ratio α (c); and average baroclinic vector (d) for summer temperature extremes (cold in blue, warm in red, all data in black) over Eastern North America as a function of the extreme co-percentile q. The values are expressed as average deviations from the respective seasonal cycles. The error bars indicate one standard deviation of the mean and are only shown for the continuous lines.

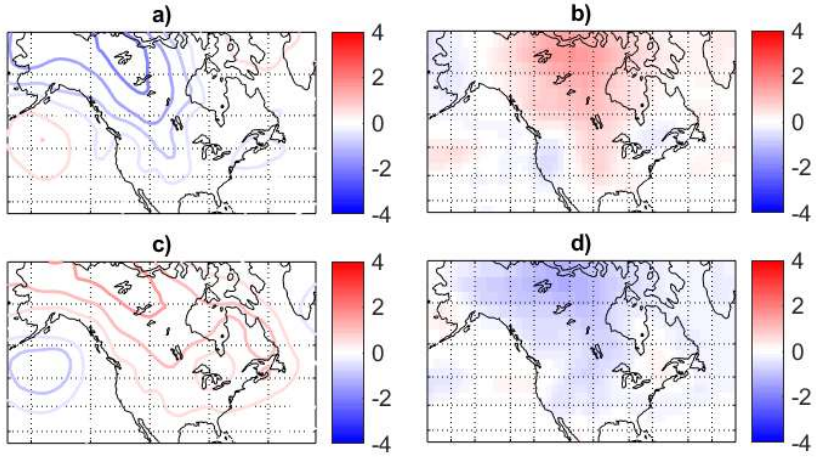

Fig. 10 Composite sea-level pressure (SLP) anomalies in hPa (a,c) and temperature anomalies (T2M) in K (b,d) corresponding to the 10% highest (a,b) and lowest (c,d) values of deviations from the seasonal cycle of the co-recurrence ratio α during JJA. There are no regions where at least 2/3 of the composited anomalies have the same sign.

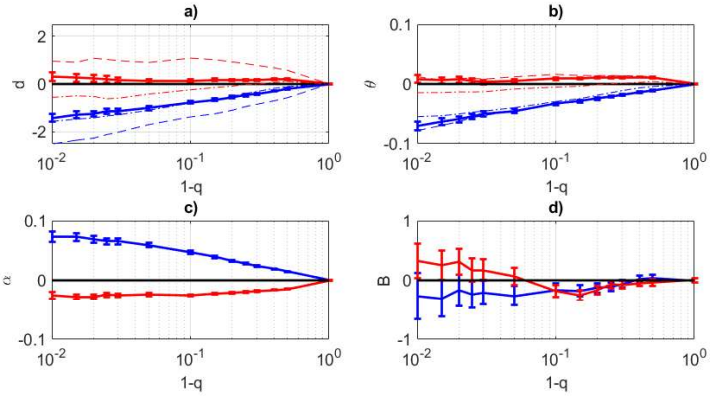

Fig. 11 Local dimensions d_{SLP} (dashed), d_{T2M} (dashed-dotted) and co-dimension $d_{SLP,T2M}$ (continuous) (a); inverse persistences θ_{SLP} (dashed), θ_{T2M} (dahsed-dotted) and inverse co-persistence $\theta_{SLP,T2M}$(continuous) (b); co-recurrence ratio α (c); and average baroclinc vector (d), for winter temperature extremes (cold in blue, warm in red, all data in black) over Eastern North America as a function of the extreme co-percentile q. The values are expressed as deviations from the respective seasonal cycles. The error bars indicate one standard deviation of the mean and are only shown for the continuous lines.

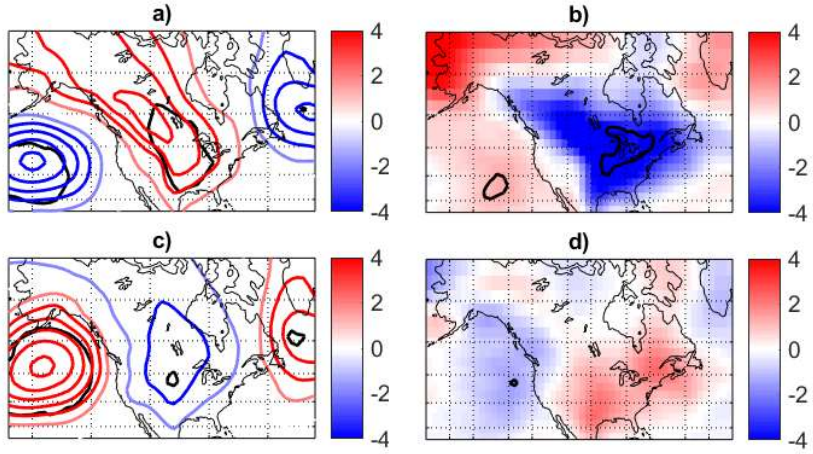

Fig. 12 Composite sea-level pressure (SLP) anomalies in hPa (a,c) and temperature anomalies (T2M) in K (b,d) corresponding to the 10% highest (a,b) and lowest (c,d) values of deviations from the seasonal cycle of the co-recurrence ratio α, during DJF. The black lines bound regions where at least 2/3 of the composited anomalies have the same sign.

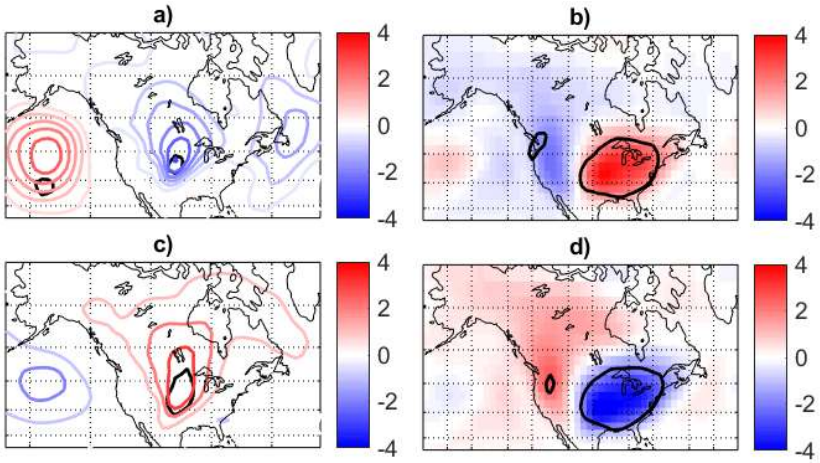

Fig. 13 Composite sea-level pressure (SLP) anomalies in hPa (a,c) and temperature anomalies (T2M) in K (b,d) corresponding to the 10% coldest (a,b) and warmest (c,d) days during JJA. The values are expressed as deviations from the respective seasonal cycles. The black lines bound regions where at least 2/3 of the composited anomalies have the same sign.

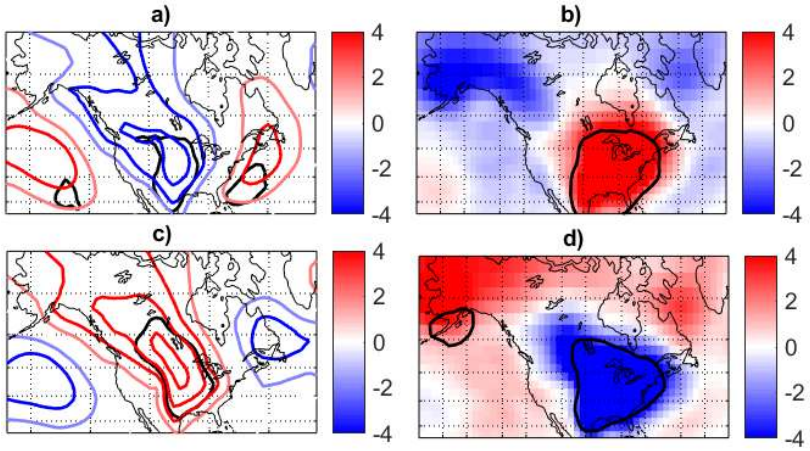

Fig. 14 Composite sea-level pressure (SLP) anomalies in hPa (a,c) and temperature anomalies (T2M) in K (b,d) corresponding to the 10% coldest (a,b) and warmest (c,d) days during DJF. The values are expressed as deviations from the respective seasonal cycles. The black lines bound regions where at least 2/3 of composited anomalies have the same sign.