



Detecting hotspots of atmosphere–vegetation interaction via slowing down – Part 1: A stochastic approach

S. Bathiany¹, M. Claussen^{1,2}, and K. Fraedrich^{1,2}

¹Max Planck Institute for Meteorology, KlimaCampus Hamburg, Germany

²Meteorologisches Institut, Universität Hamburg, KlimaCampus Hamburg, Germany

Correspondence to: S. Bathiany (sebastian.bathiany@zmaw.de)

Received: 27 June 2012 – Published in Earth Syst. Dynam. Discuss.: 20 July 2012

Revised: 18 September 2012 – Accepted: 30 October 2012 – Published: 26 February 2013

Abstract. An analysis of so-called early warning signals (EWS) is proposed to identify the spatial origin of a sudden transition that results from a loss in stability of a current state. EWS, such as rising variance and autocorrelation, can be indicators of an increased relaxation time (slowing down). One particular problem of EWS-based predictions is the requirement of sufficiently long time series. Spatial EWS have been suggested to alleviate this problem by combining different observations from the same time. However, the benefit of EWS has only been shown in idealised systems of predefined spatial extent. In a more general context like a complex climate system model, the critical subsystem that exhibits a loss in stability (hotspot) and the critical mode of the transition may be unknown.

In this study we document this problem with a simple stochastic model of atmosphere–vegetation interaction where EWS at individual grid cells are not always detectable before a vegetation collapse as the local loss in stability can be small. However, we suggest that EWS can be applied as a diagnostic tool to find the hotspot of a sudden transition and to distinguish this hotspot from regions experiencing an induced tipping. For this purpose we present a scheme which identifies a hotspot as a certain combination of grid cells which maximise an EWS. The method can provide information on the causality of sudden transitions and may help to improve the knowledge on the susceptibility of climate models and other systems.

1 Introduction

The existence of potential tipping points in the climate system has received growing attention in recent years (Lenton et al., 2008; Lenton, 2011). In the narrower sense, a tipping point occurs when a system becomes very susceptible to an external forcing due to large positive feedbacks. In the extreme case the system's attractor disappears at a threshold value of the forcing (bifurcation) and the state has to approach a different attractor.

In order to predict the collapse at a preconceived bifurcation or to distinguish such changes in stability from random state transitions, it has been proposed to exploit statistical precursors of instabilities (Wiesenfeld, 1985a,b, Wiesenfeld and McNamara, 1986), also called early warning signals (EWS; Scheffer et al., 2009). The fundamental assumption behind their applicability is that the system is close to a deterministic solution and perturbed by small fluctuations which can be described as white noise. In case of the climate system this approach resembles Hasselmann's concept of stochastic climate models (Hasselmann, 1976). When the system's stable fixed point loses stability when approaching a local bifurcation (e.g. a saddle-node bifurcation), an eigenvalue approaches 0 (if time is continuous). As a result, the linear relaxation time of the corresponding mode increases (Wissel, 1984). This phenomenon has recently been referred to as "critical slowing down" (Rietkerk et al., 1996; Scheffer et al., 2009; Ditlevsen and Johnsen, 2010; Dakos et al., 2010, 2011; Lenton, 2011; Lenton et al., 2012b). To avoid confusion with the phenomenon of algebraic (rather than exponential) decay in systems with second-order phase transitions (Strogatz, 1994), we will refer to the increased relaxation time simply

as “slowing down”. As a consequence of slowing down, the system’s autocorrelation and variance can increase (Scheffer et al., 2009), and the spectrum is reddened (Kleinen et al., 2003). Considering non-linear terms in the stability analysis, it follows that the skewness of the state variable can also increase in magnitude (Guttal and Jayaprakash, 2008).

However, the external parameter change must be slow enough for the system to stay close to equilibrium and to allow sufficiently long time series for a statistically significant detection of EWS. A lack of detectability can thus impede any final conclusion on the existence of slowing down prior to an abrupt event. For example, Dakos et al. (2008) detected a consistent increase in autocorrelation with 95 % probability in only 2 out of 8 palaeo records (see their Table S3), and the results seem to depend on the choice of the analysis method, parameter values and the particular proxy (Lenton et al., 2012a,b). This problem becomes worse close to the tipping point (for example see Dakos et al., 2012) because the uncertainty of an estimate from one sample of a fixed number of data points increases. In statistical terms, the sampling variances of the estimators of variance and autocorrelation increase with autocorrelation (Priestley, 1981). Better resolved time series may not always provide a solution as a sampling below the dynamic timescale of the system will not add relevant information.

Instead, the use of spatial EWS has been suggested (Guttal and Jayaprakash, 2009; Donangelo et al., 2010; Dakos et al., 2010): in analogy to the time domain, spatial variance and cross-correlations between different units of a spatially explicit system as well as the spatial correlation length increase towards a tipping point. Such spatial EWS use each time step as a sample to infer the stability, while temporal EWS need a window of many subsequent time steps. As forcing changes over time in transient cases, temporal EWS thus involve information on previous states of the system. It is therefore often argued that spatial EWS could provide a more precise estimate of the current stability. However, in these previous studies on spatial EWS, the system’s boundaries are known and well-defined. In addition, the application of the one-dimensional concept of EWS to high-dimensional systems, though justified by theory (Ditlevsen and Johnsen, 2010; Sieber and Thompson, 2012), in practice requires a priori knowledge on the critical mode of the transition (Held and Kleinen, 2004). This critical mode indicates in which direction in phase space the bifurcation occurs and thus how the information should be combined to yield EWS.

In this study, we consider the case where both, the critical mode as well as the critical subsystem, are unknown. First, we demonstrate that under such general conditions EWS may not be detectable at any particular location of the system.

Second, we propose an alternative application of EWS: the diagnostic detection of critical regions of slowing down (hotspots) in time series.

The potential tipping point we analyse is the decline of North African vegetation cover in the mid-Holocene. In the




Fig. 1. Stability diagram for the one-dimensional conceptual model with $k = 300$. Blue lines: equilibrium precipitation P^* as calculated from $P^*(V) = P_d + kV$ for different P_d . Green line: equilibrium vegetation cover $V^*(P)$ (Eq. 1).

Sahara and Sahel region, vegetation cover and precipitation are considered to be linked by a positive feedback on timescales beyond years (Claussen, 2009). The reasons are the effect of surface albedo on atmospheric stability (Charney, 1975), and the vegetation’s contribution to water recycling (Claussen, 1997; Hales et al., 2004). In models with a large atmosphere–vegetation feedback, two stable equilibria can exist (Claussen, 1998; Brovkin et al., 1998; Zeng and Neelin, 2000; Wang and Eltahir, 2000; Irizarry-Ortiz et al., 2003) and the gradual change in orbital forcing can cause a sudden collapse in vegetation cover (Claussen et al., 1999; Liu et al., 2006).

Our study is structured as follows: in Sect. 2 we present a stochastic model of atmosphere–vegetation interaction which produces a vegetation collapse when a control parameter is varied. We then use the stochastic model to document a specific limitation of local EWS in a spatially explicit setting (Sect. 3). Based on this finding, we explain our concept of a hotspot and present an algorithm for the detection of hotspots of slowing down (Sect. 4). We then discuss the performance of this algorithm for different properties of the analysed time series and different parameter choices and conclude in Sect. 5 by discussing possible applications and limitations of our approach. An application of our method to the results of an atmosphere–vegetation model of intermediate complexity will be presented in a subsequent article.

2 A stochastic model of atmosphere–vegetation interaction

In order to test the performance of EWS-related methods, we generate time series with a simple stochastic model of

vegetation dynamics in subtropical deserts. The structure of this model is similar to the conceptual model of Brovkin et al. (1998), Wang (2004), and Liu et al. (2006): annual precipitation P is a linear function of vegetation cover V , while equilibrium vegetation cover V^* as a function of P is of sigmoidal shape (Fig. 1):

$$V^* = \begin{cases} 0 & \text{if } P < P_1 \\ 1 & \text{if } P > P_2 \\ 1.03 - \frac{1.03}{1 + \alpha \left(\frac{P - P_1}{\exp(\gamma \delta)} \right)^2} & \text{otherwise,} \end{cases} \quad (1)$$

with

$$P_1 = \beta \exp(\gamma \delta/2) \\ P_2 = \beta \exp(\gamma \delta/2) + \frac{\exp(\gamma \delta)}{\sqrt{0.03 \alpha}}.$$

This function is the result of a semi-empirical fit to observations (Brovkin et al., 2002) and referred to as the original VE-CODE model in Bathiany et al. (2012). Parameter values in all our simulations are $\alpha = 0.0011$, $\beta = 28$, $\gamma = 1.7 \times 10^{-4}$, and $\delta = 9100$. For all time series we analyse in this study, P is always between P_1 and P_2 .

If the conditions of a specific region are described with only one value of each, V and P , the system's deterministic equilibria can be depicted as intersections of the green and blue curve in Fig. 1. Reducing the external parameter P_d describes the effect of decreasing Northern Hemisphere summer insolation during the mid-Holocene, leading to a decrease in precipitation. When the green equilibrium disappears the system experiences a saddle-node bifurcation and vegetation cover has to collapse to the remaining desert state.

We extend this conceptual model by defining V and P for several elements with index i (for example to represent different grid cells in a climate model). At each of the N elements, equilibrium vegetation cover depends only on the local precipitation according to $V^*(P)$. Vegetation cover is updated every timestep via the dynamic equation

$$V_i^{t+1} = V_i^t + \frac{V^*(P_i^t) - V_i^t}{\tau} \Delta t + \sigma_V \sqrt{\Delta t} \eta_i^t. \quad (2)$$

As P represents mm yr^{-1} our time step is 1 yr, so $\Delta t = 1$. The timescale τ describes how fast vegetation cover can establish in previously unvegetated areas (or die back in vegetated areas). Following Liu et al. (2006) we fix τ to 5 yr which is meant to represent the dynamics of grass in arid subtropical ecosystems. Due to atmospheric water transport and circulation changes, local precipitation P_i at a particular time t depends on vegetation cover at all elements:

$$P_i^t = \underbrace{P_{0i} + s_i B^t}_{P_d} + \underbrace{\sum_{j=1}^N \mathbf{k}_{ij} V_j^t}_{P^*} + \sigma_P \eta_i^t. \quad (3)$$

Due to the fast equilibration time of the atmosphere, Eq. (3) is not dynamic, and the V_i are all the state variables of this dynamical system. The system is globally coupled via \mathbf{k} and in this regard differs from reaction-diffusion models with interactions between adjacent elements only. The choice of $V^*(P)$ and the interaction matrix \mathbf{k} determine the strength and spatial structure of the atmosphere–vegetation feedback and thus the stability properties of the system.

Brovkin et al. (1998), Wang (2004), and Liu et al. (2006) refer to the equilibrium precipitation in the absence of any vegetation as P_d . However, as P_d may differ at different elements, we split it into P_{0i} , which is variable in space but not in time, and $s_i B$ with a scalar B as external control parameter. The local sensitivity of background precipitation to B is determined by parameters s_i , which are also variable in space, but not in time. Hence, in all systems analysed in this article, B is the bifurcation parameter, determined by one single number. In physical terms, B describes the effect of climate forcings, while the numbers we use are chosen arbitrarily. Also following Brovkin et al. (1998) and Wang (2004) we call $P^*(V)$ the equilibrium precipitation at a particular location (Fig. 1). P^* can be interpreted as precipitation in the noise-free case or as the long-time mean when vegetation cover is fixed at a permanent value.

The Gaussian white noise process η with zero mean and small noise level σ is uncorrelated in space. We distinguish two types of noise but always use only one of them in our experiments: σ_V controls perturbations which are added to Eq. (2) directly (additive noise), while σ_P controls perturbations added to precipitation and whose impact on the state variable V_i depends on the system's state itself (multiplicative noise). Atmospheric variability is more realistically accounted for by the multiplicative noise case, whereas the additive noise case may describe disturbances other than atmospheric conditions, such as fire, diseases or grazing. Only the additive noise case allows rising variance to be a generic indicator of slowing down (Dakos et al., 2012), although we will show that in our specific model rising variance is also a useful indicator in the multiplicative noise case. In all our simulations we use very small noise levels of $\sigma_V = 0.00013$ or $\sigma_P = 2$. For simplicity, we provide P_d , k_{ij} and σ_P without units, although the value of P represents mm yr^{-1} .

3 Performance of early warning signals (EWS) in spatially coupled systems

In the following, we address the limitations of EWS at individual elements. All statistical indicators are calculated from time series of the state variables V_i . Autocorrelations are determined for lag 1, cross-correlations for lag 0.




Fig. 2. Characteristics of system 1 in dependency on parameter B . **(a)** Equilibrium vegetation cover, **(b)** autocorrelation (lag 1), **(c)** variance (additive noise only), **(d)** variance (multiplicative noise only).

3.1 First example: induced tipping

Consider the following simple system (system 1): 2 elements are coupled in a way that the first element can be bistable due to a large local feedback between P and V . Precipitation at the second element depends on vegetation cover at the first element, but not vice versa. We implement this property by choosing the interaction matrix

$$\mathbf{k} = \begin{pmatrix} 300 & 0 \\ 200 & 0 \end{pmatrix}$$

and parameters

$$P_0 = \begin{pmatrix} 0 \\ 100 \end{pmatrix}$$

$$s = \begin{pmatrix} 1 \\ 0.1 \end{pmatrix}.$$

As B is reduced, element 2 (blue) collapses in response to the collapse of element 1 (red; Fig. 2a). The curves in all our figures are derived from stationary time series. However, if B was reduced very slowly during an experiment, the transient time series of the collapse would follow the equilibrium curves, i.e. in Fig. 2a, very closely because the noise level is small and because the timescales of both elements are identical and small compared to the parameter change. Therefore, it would not be possible to infer the causality of a transition from the timing of the collapses at different elements.

As the collapse of element 2 is induced by element 1 it is not related to a substantial loss of its own stability. It rather experiences the transition as an induced tipping caused by a sudden change in external conditions that are imposed by element 1. The stability of element 2 is hardly affected by B directly as the difference in s_1 and s_2 indicates.

Therefore, element 1 shows a clear increase in autocorrelation (Fig. 2b) and variance (Fig. 2c) in the additive noise case,




Fig. 3. Characteristics of system 2 in dependency on parameter B for versions with a different number of elements. **(a)** Equilibrium vegetation cover (identical for any number of elements), **(b)** autocorrelation (lag 1), **(c)** cross-correlation (no lag), **(d)** variance (additive noise only), **(e)** variance (multiplicative noise only). Note that all elements of a specific system are identical and thus have the same measured indicators.

but element 2 does not. Only when the noise is multiplicative the system under consideration shows an increased variance (Fig. 2d; note that the scale differs from Fig. 2c by a factor 100), but results for autocorrelation are similar to the additive noise case. The increase in variance in the multiplicative noise case is specific to the conceptual model and results from the increasing sensitivity of V^* to precipitation changes when P is reduced (Fig. 1). In case of a single isolated element without any P - V -feedback ($\mathbf{k} = 0$) there would still be an increase in variance in the multiplicative noise case, but not in the additive noise case. In our system 1, the slowing down at element 1 also affects element 2 due to the interaction term. This is the reason for the rise of the blue curve in Fig. 2c.

To obtain sufficiently precise estimates of the statistical properties in Fig. 2, we performed stationary time series of 10 million data points each for different values of B . In a

Table 1. Parameters P_{0i} and s_i in example system 3 for 4 different types of elements. Colours correspond to those in Fig. 4.

	red	blue	green	brown
P_0	−50	40	210	40
s	1.7	0.8	0.2	0.9

transient situation where the sampling error is much larger, the collapse of element 2 would hardly be predictable with EWS.

3.2 Second example: collective bistability

To pursue this further, we now consider a system (system 2) with a different number of elements, distinguishing versions with 1, 2, 5, 10, and 20 elements, where any particular element has the identical parameters $P_{0i} = 0$, $s_i = 1$, and $k_{ij} = 300/N$. By dividing the entries of interaction matrix \mathbf{k} by the number of elements in the system, we equally distribute the P - V -feedback over all elements. When more and more elements are coupled, the spatial resolution increases but the bifurcation diagram of this globally coupled system (Fig. 3a) does not change. As local feedbacks (determined by k_{ij}) are weak, no single element would be bistable anymore if all other elements were fixed. This fact distinguishes our model from those in Guttal and Jayaprakash (2009), Dakos et al. (2010) and Donangelo et al. (2010), where individually bistable elements are coupled. However, the system as a whole still shows a bifurcation due to the spatial interactions k_{ij} with $i \neq j$.

As we couple more and more elements, it is evident that EWS like rising autocorrelation and variance at individual elements, as well as rising cross-correlation, tend to disappear (Fig. 3b–d). Again, variance in the multiplicative noise case (Fig. 3e) is an exception due to the increased slope in $V^*(P)$.

The one element-case here (red curves in Fig. 3) is identical to element 1 from system 1 (red curves in Fig. 2), and also to the system in Fig. 1 in Bathiany et al. (2012). For EWS to appear exactly like in this single element case, the system's time series need to be projected on the critical mode of the transition, a technique introduced as “degenerate fingerprinting” by Held and Kleinen (2004). The critical mode implies the direction in phase space in which the bifurcation occurs. Slowing down particularly occurs for this mode and can be revealed by the appropriate projection. In contrast, other modes of the system's variability are not necessarily influenced by slowing down as the changes of the stability landscape in other directions (characterised by changes of the according eigenvalues) are unrelated to the bifurcation. Hence, EWS in projections on other modes cannot be expected. The analysis of local EWS at individual elements would generally contain information on these other modes of variability and would therefore be a futile strategy. It has to

Table 2. Interaction matrix \mathbf{k} in system 3, distinguishing 4 different types of elements. Colours correspond to those in Fig. 4. A number in some row A and column B stands for the impact of any single element of type B on any single element of type A (for example: impact of red on blue: 15, impact of blue on red: 5).

	red	blue	green	brown
red	27	5	10	10
blue	15	4	3	3
green	8	2	15	2
brown	2	3	2	5

be concluded that if the critical mode of the transition is not known beforehand, the tipping can be unpredictable even in cases of very long time series.

4 Early warning signal – based hotspot detection method

So far we have chosen systems of simple structure. In a more general case like a spatially resolved climate model, the stability structure will be more complicated. Certain subsystems of the climate may show a loss of stability during a change in external forcing while the rest of the system may respond only indirectly, or even evolve independently. In Sect. 3 we documented that in multidimensional settings individual elements can fail to show EWS before a sudden transition. While this constitutes a caveat for the prediction of sudden transitions, one may make a virtue out of this caveat by using EWS to diagnostically infer information on the causality of a sudden transition. In terms of system 1, we aim at finding the nucleus of slowing down (hotspot) by distinguishing elements of the red and the blue kind. This is not possible by looking at the system's state directly because red and blue elements collapse in synchrony. Of course, in complex systems there will be a continuum from red to blue and the definition of a threshold in between will be somewhat arbitrary. In principle, we expect that the hotspot can be identified as the combination of elements which (when projected on their critical mode) maximises an indicator of slowing down. In the following, we present an algorithm for hotspot detection which we apply to our stochastic model.

4.1 Highly idealised North African vegetation dynamics

As yet another example of the stochastic model framework in Sect. 2, consider 25 elements which can be interpreted as a highly idealised Northern Africa (Fig. 4). We refer to this system as system 3. Again we choose parameter values which lead to preconceived properties of the model: 5 of the 25 elements gradually become desert when B is reduced (brown elements). 5 elements stay mostly vegetated (green elements), a set of 9 elements becomes bistable and finally

Table 3. Example signal list for elements 19, 20 and 25 in system 3. Parameter settings correspond to set 1 in Table 4.

Area	Signal $\times 1000$
19	9.1611
20	4.8099
19, 20	11.5094
25	1.0391
19, 25	7.8746
20, 25	4.0192
19, 20, 25	11.7716
Weights (19,20,25):	40.32, 32.11, 24.70



Fig. 4. Structure of system 3. Red: area with strong P - V -feedback (hotspot), blue: passively dependent on red area, brown: dry area, green: moist area.

collapses due to a saddle-node bifurcation (red elements) and 6 elements substantially depend on the red ones but show a much weaker local atmosphere–vegetation feedback (blue elements; see Fig. 5). Elements with identical colours have identical parameter values and thus have the same state in equilibrium. Hence, there are 4 s_i and P_{0_i} (Table 1), and 16 k_{ij} (Table 2). In similarity to the examples in Sect. 3.2, no element is bistable on its own, as local feedbacks k_{ii} are too small. It is only due to the strong spatial interactions between the red elements that the system can bifurcate and thus show a vegetation collapse at $B \approx 43$.

The nucleus of the transition is the red area because this is where the system loses stability due to strong atmosphere–vegetation interaction. In the following, we refer to the red area as a hotspot.

4.2 Strategy for the detection of hotspots

We now explain our method of analysis by applying it to system 3. As several modifications of our algorithm are possible, we provide the explanation in two steps:

1. In this section we address the general strategy of our approach. This strategy sets the framework of analysis




Fig. 5. Equilibrium vegetation cover at different elements of system 3 and for different bifurcation parameter values B . The colours correspond to the elements in Fig. 4. The vertical black dashed lines indicate the values of B used for the four stationary simulations (the smallest one also lying above the tipping point). They correspond to BV1 in Table 4 and are used for Figs. 7–11 and Table 3.

which is presented in Fig. 6 and is the same for all versions of our method. To illustrate our explanation we complement our step by step description with a simple example. This example is referred to at the end of each particular step and presented in Fig. 7 and Table 3.

2. The framework of analysis presented in this section is too general to cover all technical details as presented in Fig. 7 and Table 3. These details can differ from case to case. We therefore introduce the different versions of our algorithm together with a discussion of their advantages and disadvantages in Sect. 4.3.

In all cases the analysis is applied to J preferably long stationary simulations for fixed but different forcings B_j ($j = 1, 2, \dots, J$) before the bifurcation point. In our example and for all figures which follow we choose time series of vegetation cover for $B_1 = 150$, $B_2 = 90$, $B_3 = 55$, and $B_4 = 43$ (hence $J = 4$; vertical dashed lines in Fig. 5). All steps that follow are an analysis of these time series and do not involve the model which generated them. We describe the individual steps of the analysis by starting with part B in Fig. 6, as this part of the analysis corresponds to the original degenerate fingerprinting by Held and Kleinen (2004), without time aggregation.

- B1. For a given part of the system with N_p elements, we select a subset of n elements from these N_p elements. We refer to this subset as an area. Hence, there are three levels of selected elements where each set is a subset of the previous one: the number of elements in the complete system (N ; here: 25), the number of elements in a

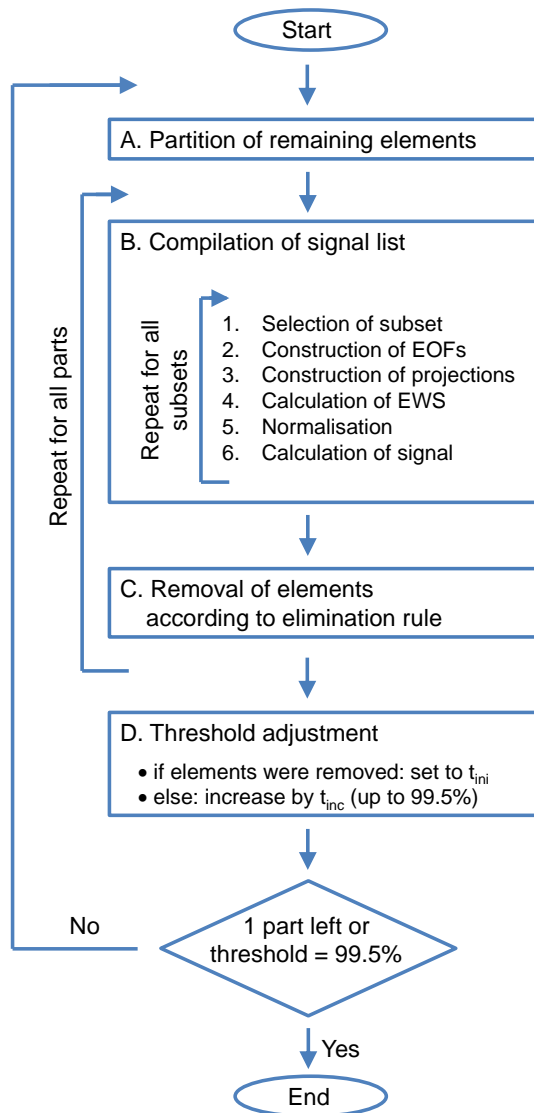


Fig. 6. General flowchart of the hotspot detection scheme.

part of the system (N_p), and the number of elements in an area of this part (n). Example: we choose elements 19, 20 and 25 as a part. Hence, $N_p = 3$, and n can be 1 (3 possible combinations), 2 (3 possible combinations) or 3 (1 possible combination).

B2. For the n selected elements, we calculate the leading empirical orthogonal function (EOF; eigenvector of the covariance or correlation matrix which represents the largest variance). To construct the EOFs, we use the freely-available linear algebra package LAPACK. In our example, there are 7 combinations of the 3 elements (top left corner of panels in Fig. 7). The 3 cases with single elements are trivial and each EOF is 1. The 3 cases with 2 elements are also trivial ($\sqrt{1/2}$, $\sqrt{1/2}$) because

the (symmetric) correlation matrix contains ones on its main diagonal.

B3. In the general case, we calculate the leading EOF for all time slices B_j from $j = 2$ to J . For every EOF $_j$, we project all time slices from B_1 to B_j on EOF $_j$. Special cases: The projection of time series B_j on EOF $_j$ is the principal component of EOF $_j$. In case of areas consisting of one single element, the projections are identical with the time series themselves. The standard version of our algorithm only involves projections on EOF $_j$ (see Sect. 4.3). In our example, there are therefore 4 projections for each area (small panels in Fig. 7).

B4. We calculate a statistical property like autocorrelation or variance of the corresponding projections to use it as an EWS. For all projections on some EOF $_j$ the result is a curve of this EWS versus B , just like those in Figs. 2b–d and 3b–e, but less well resolved (j points only). In our example, we use $J = 4$ and autocorrelation as EWS, hence there are 4 values of AC for each of the 7 areas, shown as a line plot in each panel of Fig. 7.

To automatically compare the results for different areas, we expand this degenerate fingerprinting method with the following steps:

B5 As it is not the absolute value of a statistical property but its increase which indicates slowing down, we shift the curve vertically in order to be 0 at $j = 1$. Fig. 7 (bottom right) shows the shifted AC-curves of all 7 areas of our example.

B6 Based on the aligned curves, we define a *signal* which is one number to quantify the strength of an indicator and to compare different areas. The definition of the signal can differ, but it always involves all time slices. We do so to take into account not only the difference between the first and last B , but the whole evolution of an EWS vs. B as is suggested by our results in Fig. 2. Table 3 gives an example of all areas and their associated signals for the part consisting of elements 19, 20, and 25.

We repeat steps B1–6 for all possible combinations of elements. If the N_p elements mentioned in step 1 represent the whole system under consideration ($N_p = N$), one can then determine the area with the maximum signal, or the areas with a signal above a certain threshold. However, this requires the calculation of $2^N - 1$ such signal (not 2^N because selecting 0 elements is not an option). This becomes unfeasible already for N beyond 10. Therefore, not all possible combinations can be calculated and we use an iterative selection process to decide which elements can be dropped from the analysis:

A. We randomly divide the whole system into a number of non-overlapping parts. The number of parts is calculated from the fixed parameter n_{\max} via the ceiling

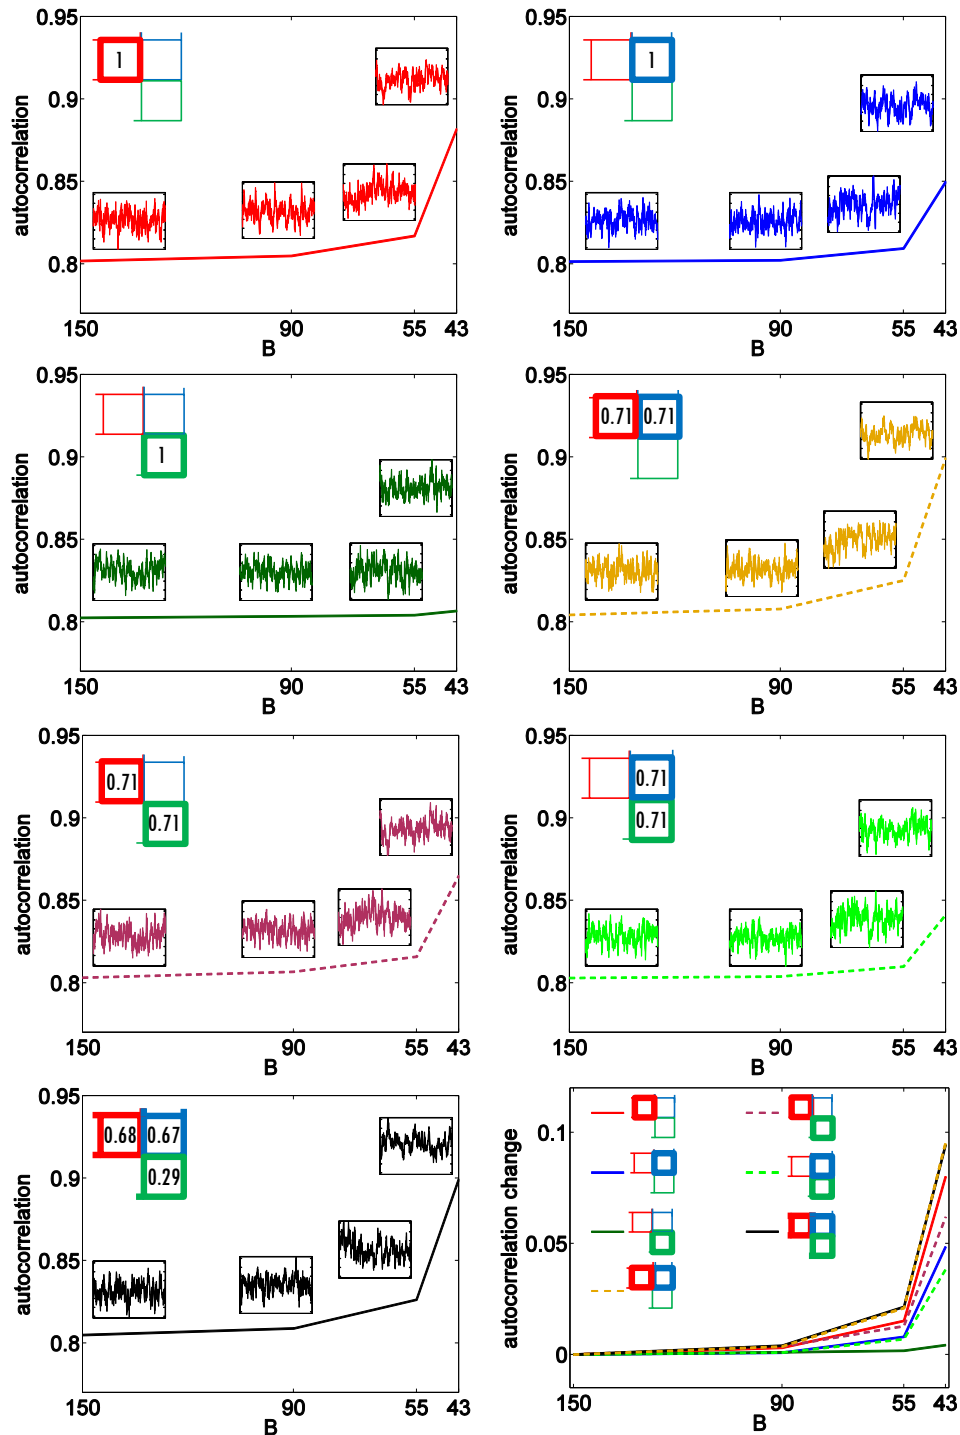


Fig. 7. Example to illustrate the hotspot detection scheme using elements 19, 20, and 25 of system 3. All panels except bottom right: in each top left corner the particular area is highlighted by bold lines. The numbers inside the elements give the eigenvectors of the correlation matrix (EOF) for $B = 43$. All four time slices (forcings BV1 as in Fig. 5) are projected on this EOF. The four time series in each panel show a chunk of 500 yr from these projections (normed to standard deviation 1 and mean 0). The autocorrelation at lag 1 for each projection is depicted as a line plot in dependency on B . Bottom right: all seven curves are shifted to 0 at $B = 150$ in order to compare the signals of the areas. Parameter settings correspond to set 1 in Table 4.

function $\lceil \frac{N}{n_{\max}} \rceil$. The number of parts is thus as small as possible for a given n_{\max} . The size of each part is then determined by distributing the N elements as equally as possible, so that each N_p fulfills $2 \leq N_p \leq n_{\max}$.

- B. For each part, steps 1–6 are applied. As an example, imagine that system 3 is analysed with $n_{\max} = 3$. Hence, the system is subdivided into 9 parts, of which 7 parts contain 3 elements, and 2 parts contain 2 elements.
- C. From a *signal list* like Table 3, the contribution of different elements can be disentangled, with the aim to drop unimportant elements from the analysis completely. The removal of an element means that it is not considered to be part of the system anymore and is not used from that point on. In this sense, the total system size N successively decreases and with it the number of parts is also reduced automatically.

In principle, our selection process resembles the logic of a football world cup: each team (element) does not compete directly against every other team, but only against those of the same group (part). Only teams performing well enough in their group remain in the tournament.

The best rules of how to remove an element can depend on the system and will be discussed under the term *elimination rule* (ER) in Sect. 4.3.

- D. As a criterion for removing elements, we set a threshold which is adjusted interactively to prevent that too many elements are removed at once and too early. The threshold is a relative number in % and relates to different things depending on the elimination rule. In our example as well as all following figures, we set the initial threshold t_{ini} to 5%. As long as no element can be removed in any part, we increase the threshold by $t_{\text{inc}} = 5\%$. If the threshold would reach or exceeded 100%, we set the threshold to 99.5%. If at least one element can be removed, we reset the threshold to its initial value t_{ini} . In both cases, we then partition the remaining elements anew, starting from step A (large loop in Fig. 6).

This way, the considered number of elements is gradually reduced. After each calculation of the signals in all possible areas of all current parts and the potential removal of elements the procedure is repeated. It ends as soon as one of the following conditions is true: (1) the total number of remaining elements is not larger than n_{\max} , in which case the analysis is repeated one last time with one part only. (2) The relative threshold reaches 99.5%, but still no elements can be removed because the remaining elements are too similar to be discriminated.

The algorithm serves as a sieve in order to filter out the important elements with a sufficiently small number of calculations. Without the removal of elements, the number of




Fig. 8. Performance of the hotspot detection algorithm for system 3 with additive noise using time series of 2000 yr. The frequencies show the number of times a particular element remains until the end of the selection process for 500 repetitions. Each repetition involves the generation of a new time series and its analysis with the hotspot detection algorithm. The solid black line marks the expectation value for a random selection where all elements are selected with equal probability. The red dashed line marks the 95% probability threshold of the corresponding cumulative binomial distribution. Parameter settings correspond to set 1 in Table 4.

possible combinations would be too large to achieve a robust hotspot detection within a feasible amount of time. As the results depend on the random distribution of elements to different parts, they will be very similar but not completely identical when the analysis is repeated. The hotspot of slowing down can be identified if the time series are long enough (or if enough realisations are available) because the remaining elements at the end of the analysis tend to contribute most to slowing down.

To obtain more quantitative results, all signals calculated during the procedure can be collected in a sorted list for further analysis. Elements belonging to the hotspot tend to be part of the areas with the strongest signals and are on top of the list. However, elements that have been removed early during the analysis are not well sampled. The method therefore only provides information on the nature of the hotspot, but less on the stability properties of the rest of the system.

4.3 Parameter options and performance analysis

It has become obvious in the previous sections that the algorithm involves a number of options and parameter values which have to be chosen in advance. Also, the performance of the method will depend on properties of the original time series. For a quantitative comparison of the algorithm's performance under different conditions (parameter settings,

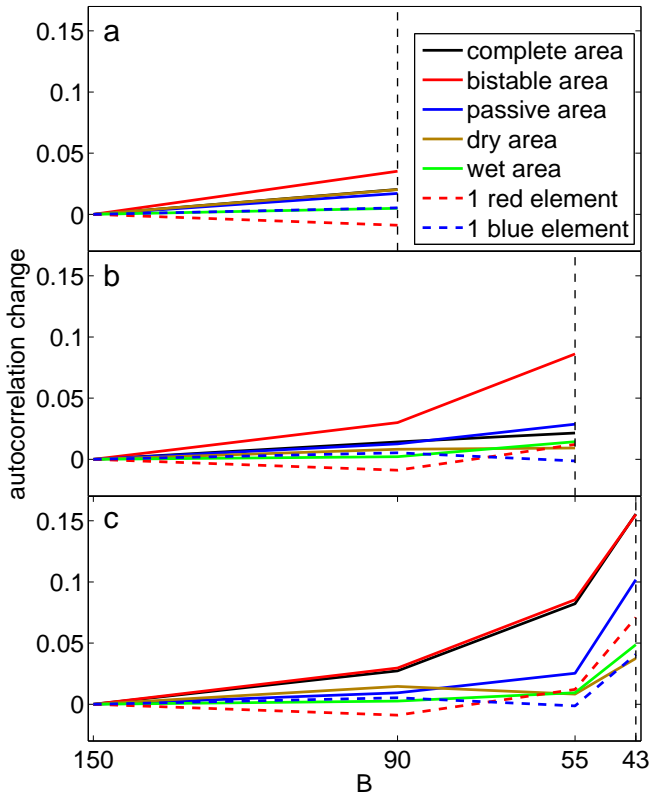


Fig. 9. Autocorrelation changes of projections on leading EOFs for system 3 with multiplicative noise. The leading EOFs have been calculated for (a) $B_2 = 90$, (b) $B_3 = 55$, (c) $B_4 = 43$. In each case, all previous time series (including the one used for the EOF) are projected on the according EOF. The analysis is applied to the full system (black) as well as only parts of the system (other colours). The colours correspond to the elements in Fig. 4. Parameter settings correspond to set 21 in Table 4.

choice of algorithm and time series properties), we perform 500 Monte Carlo experiments for each condition (Table 4).

In each experiment a new realisation of the time series is generated with system 3 and then analysed with the hotspot detection algorithm. Figure 8 shows how often each element remains until the end of each experiment for the additive noise case and a time series length of 2000 yr. After the 500 repetitions, we evaluate which fraction f_1 of the $500 \times n_{\max}$ potentially identified elements belongs to the hotspot, and which fraction f_2 of the actually obtained elements belongs to the hotspot. f_1 and f_2 can differ because it is not always n_{\max} elements that remain in the end.

As a measure of the method's performance η , we define for both variants of f :

$$\eta_{1,2} = (f_{1,2} - \frac{H}{N}) / (1 - \frac{H}{N}), \quad (4)$$

with N as the size of the system (25) and H as the size of the hotspot (9). If we assume that all 25 elements have an equal chance to be selected, the probability for any ob-

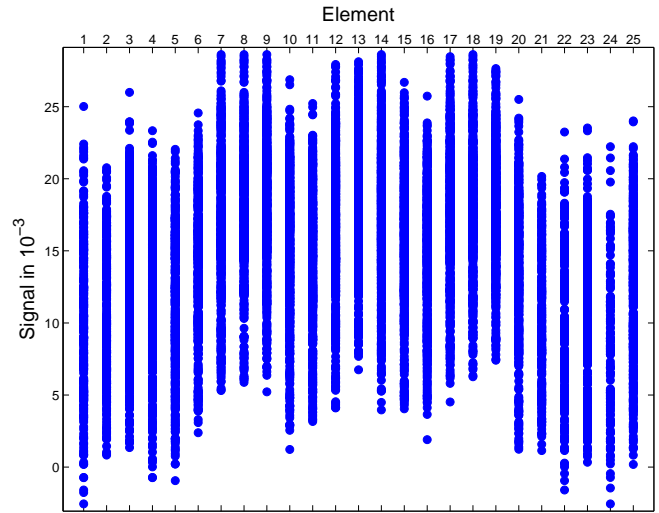


Fig. 10. Signal list for system 3 with additive noise using time series of 100 000 yr. Ordinate: absolute signal; abscissa: elements of the system. Any area that has been calculated during the analysis is represented at the ordinate value of its signal. All elements that are part of this area are marked as blue dots. Parameter settings correspond to set 1 in Table 4.




Fig. 11. Signal list for system 3 with multiplicative noise using time series of 10 000 yr. Ordinate: absolute signal; abscissa: elements of the system. Any area that has been calculated during the analysis is represented at the ordinate value of its signal. All elements that are part of this area are marked as blue dots. Parameter settings correspond to set 21 in Table 4.

tained element to be part of the hotspot is $H/N = 9/25$. A detection which does not differ from this random case has performance 0.

If exactly n_{\max} elements are returned in every experiment, a detection which only returns hotspot elements has performance 1 for both variants of f (which is of course only possible because we choose an n_{\max} smaller than the hotspot).

The expectation value for the occurrence of every element would be 100 in case of performance 0 (the solid black line in Fig. 8), and $500 \times 5/9$ in case of performance 1 (end of vertical scale in Fig. 8). Potential deviations from n_{\max} elements in the output can lead to performances lower than 0 and larger than 1 if we apply f_1 .

The decision for 500 repetitions can be justified by bootstrapping our Monte Carlo results (Efron, 1979): for any list of 500 sets of residual elements we draw n sets and measure their performances. We calculate the standard deviation of the obtained performances for many different n . It turns out that for 500 repetitions, the standard deviation is approx. 0.015 and rather independent of the parameter and time series properties. Therefore, we round all performances in Table 4 two decimal places. Above 500 repetitions, the uncertainty of the performance decreases very slowly while the computation time for the Monte Carlo experiments increases beyond feasibility.

From theoretical considerations, the performances in Table 4 as well as the qualitative appearance of the resulting signal lists, we draw the following conclusions with regard to different parameter choices and time series properties:

- Different *EWS* can be used within the same framework. Here, we use the increase in autocorrelation (AC) and the relative increase in variance (var). Relative increases in variance usually show better performances because of a larger signal-to-noise ratio, in agreement with Ditlevsen and Johnsen (2010). However, AC is the more generic *EWS* and also works if multiplicative noise leads to a reduction in variance (Dakos et al., 2012). For Figs. 7–11 and Table 3 we use AC as an *EWS*.
- We distinguish two *signal definitions* (SD). The most simple approach is to only use the projections on the EOF of the last time slice (B_4 in our example). The analysis is based on the assumption that this pattern resembles the critical mode, if the selected area is the hotspot. We then integrate the *EWS*-curve over B (calculate the area of the $J - 1$ segments). We refer to this signal definition as SD1. SD1 is used in our example Fig. 7 and Table 3.

An alternative (referred to as SD2) is to also consider projections on previous EOFs. This approach can add information if the leading EOF smoothly approaches the critical mode when B approaches the Tipping Point. We thus obtain $J - 1$ curves of an *EWS* vs. B (Fig. 9). To calculate the signal for a specific area, we perform a double integration. In terms of Fig. 9: first, we calculate the area under a curve with a certain colour for EOF $_{B=90}$ (Fig. 9a), EOF $_{B=55}$ (Fig. 9b), and EOF $_{B=43}$ (Fig. 9c). The resulting trajectory of integrated *EWS* is then again integrated over B . This way, not only the

shape of the projection on the last EOF is accounted for, but also the shape of previous projections.

- The choice of an *elimination rule* (ER) should be adapted to the signal definition. Again, we distinguish two elimination rules, ER1 and ER2. For SD1 we use ER1, which works as follows: for each specific element, we add up the signals of all areas this element is part of (last row, second column in Table 3), and refer to it as the element's *weight*. The threshold to remove unimportant elements is defined relative to the maximum weight of all elements in a specific part. The absolute value of the threshold therefore depends on the maximum weight and differs among the system's parts, while the relative threshold is a parameter that is independent of the parts. For example, a threshold of 70 % means that all elements with a weight smaller than 70 % of the maximum weight are removed. In our example (Fig. 7 and Table 3), element 19 belongs to the hotspot, so it contributes more to the signal than elements 20 and 25, whose weight is therefore smaller. Element 25 has a particularly small weight (24.70) as it neither belongs to the hotspot nor is it much affected by it. Its relative weight compared to the maximum weight of 40.32 is below 70 %. It would therefore be removed from the analysis if the threshold is above 70 %.

For SD2 we use a simpler approach, referred to as ER2: we divide the signal list in the set of signals above the current threshold and the set of signals below this threshold. All elements which are part of any area above the threshold remain, the other elements are removed. Hence, the threshold is directly applied to the signals itself without the calculation of weights. This measure allows a better discrimination of the elements. In the additive noise case it cannot be applied because there the maximum signal usually belongs to the complete area.

- *EOFs* can be calculated as an eigenvector of the system's covariance matrix or alternatively its correlation matrix. If based on the covariance matrix, elements with large variance will be emphasised. Whether this improves the performance of a hotspot detection generally depends on the system under analysis. In case of system 3 with multiplicative noise, variance is enhanced particularly at the hotspot when B approaches the bifurcation point. Therefore, SD2 with ER2 yield the most significant results when using covariance-based EOFs.

In general, other signal definitions and elimination rules could be devised that may be tailored to a specific system.

The most generic approach is to use SD1 in combination with ER1, correlation-based EOFs and autocorrelation as *EWS*. The last EOF (the only one used in SD1) must resemble the critical mode in any system approaching a bifurcation (although it can be difficult to come very close to the

Table 4. Performances of the hotspot detection algorithm for different sets of versions, parameter choices and time series properties. SD=signal definition, ER=elimination rule, BV=vector of forcings B_j . Performances are calculated from fractions f_1 , italic results in parenthesis from f_2 (Sect. 4.3).

Set	Parameters for hotspot detection							Time series properties						
	EWS	SD	ER	EOF	n_{\max}	t_{ini}	t_{inc}	BV	τ	noise	$T = 1000$	$T = 2000$	$T = 5000$	$T = 10000$
1	AC	1	1	corr.	5	5 %	5 %	BV1	5	add.	0.16 (0.22)	0.27 (0.33)	0.41 (0.50)	0.56 (0.69)
2	AC	1	1	covar.	5	5 %	5 %	BV1	5	add.	0.13 (0.19)	0.24 (0.30)	0.43 (0.54)	0.54 (0.68)
3	AC	1	1	corr.	5	5 %	5 %	BV1	2.5	add.	0.29 (0.36)	0.40 (0.48)	0.55 (0.70)	0.66 (0.84)
4	AC	1	1	corr.	3	5 %	5 %	BV1	5	add.	0.19 (0.22)	0.29 (0.34)	0.44 (0.51)	0.58 (0.67)
5	AC	1	1	corr.	7	5 %	5 %	BV1	5	add.	0.13 (0.18)	0.24 (0.31)	0.39 (0.49)	0.49 (0.60)
6	AC	1	1	corr.	5	5 %	5 %	BV2	5	add.	0.10 (0.13)	0.13 (0.18)	0.23 (0.30)	0.37 (0.46)
7	AC	1	1	corr.	5	5 %	5 %	BV3	5	add.	0.04 (0.08)	0.12 (0.16)	0.21 (0.27)	0.36 (0.43)
8	AC	1	1	corr.	5	5 %	1 %	BV1	5	add.	0.18 (0.19)	0.28 (0.30)	0.45 (0.47)	0.63 (0.67)
9	AC	1	1	corr.	5	5 %	2,5 %	BV1	5	add.	0.17 (0.20)	0.28 (0.32)	0.46 (0.51)	0.62 (0.69)
10	AC	1	1	corr.	5	5 %	7,5 %	BV1	5	add.	0.17 (0.24)	0.22 (0.31)	0.40 (0.51)	0.49 (0.68)
11	AC	1	1	corr.	5	5 %	10 %	BV1	5	add.	0.13 (0.23)	0.20 (0.31)	0.37 (0.53)	0.47 (0.70)
12	AC	1	1	corr.	5	5 %	12,5 %	BV1	5	add.	0.12 (0.23)	0.21 (0.35)	0.32 (0.52)	0.44 (0.67)
13	AC	1	1	corr.	5	5 %	15 %	BV1	5	add.	0.08 (0.23)	0.21 (0.37)	0.28 (0.56)	0.37 (0.70)
14	AC	1	1	corr.	5	5 %	17,5 %	BV1	5	add.	0.08 (0.24)	0.17 (0.37)	0.30 (0.54)	0.41 (0.67)
15	AC	1	1	corr.	5	5 %	20 %	BV1	5	add.	0.11 (0.27)	0.19 (0.37)	0.30 (0.54)	0.27 (0.69)
16	AC	1	1	corr.	5	5 %	30 %	BV1	5	add.	-0.01 (0.27)	0.11 (0.42)	0.26 (0.54)	0.36 (0.68)
17	AC	1	1	corr.	5	5 %	40 %	BV1	5	add.	0.11 (0.28)	0.21 (0.39)	0.28 (0.54)	0.30 (0.68)
18	AC	1	1	corr.	5	5 %	50 %	BV1	5	add.	-0.08 (0.31)	0.05 (0.39)	0.27 (0.51)	0.44 (0.59)
19	AC	1	1	corr.	5	5 %	100 %	BV1	5	add.	0.10 (0.27)	0.23 (0.33)	0.37 (0.47)	0.46 (0.54)
20	AC	1	1	corr.	5	80 %	5 %	BV1	5	add.	0.12 (0.27)	0.25 (0.39)	0.41 (0.53)	0.56 (0.69)
21	AC	2	2	covar.	5	5 %	5 %	BV1	5	mult.	0.60 (0.62)	0.66 (0.66)	0.87 (0.84)	1.10 (0.94)
22	AC	1	1	covar.	5	5 %	5 %	BV1	5	mult.	0.29 (0.36)	0.44 (0.53)	0.61 (0.71)	0.74 (0.87)
23	AC	2	2	corr.	5	5 %	5 %	BV1	5	mult.	0.43 (0.46)	0.50 (0.53)	0.70 (0.64)	0.95 (0.70)
24	AC	1	1	corr.	5	5 %	5 %	BV1	5	mult.	0.31 (0.38)	0.40 (0.48)	0.56 (0.65)	0.67 (0.81)
25	AC	2	2	covar.	5	5 %	5 %	BV2	5	mult.	0.57 (0.62)	0.64 (0.66)	0.71 (0.71)	0.82 (0.79)
26	AC	2	2	covar.	5	5 %	5 %	BV3	5	mult.	0.68 (0.69)	0.74 (0.75)	0.87 (0.85)	1.09 (0.96)
27	AC	2	2	covar.	5	5 %	5 %	BV1	2.5	mult.	0.64 (0.67)	0.79 (0.79)	1.12 (0.94)	1.63 (0.99)
28	var	1	1	corr.	5	5 %	5 %	BV1	5	add.	0.28 (0.33)	0.43 (0.51)	0.60 (0.74)	0.67 (0.86)
29	var	1	1	covar.	5	5 %	5 %	BV1	5	add.	0.29 (0.36)	0.42 (0.51)	0.62 (0.75)	0.71 (0.87)
30	var	1	1	corr.	5	5 %	5 %	BV1	2.5	add.	0.38 (0.45)	0.53 (0.63)	0.67 (0.83)	0.73 (0.94)
31	var	1	1	corr.	5	5 %	1 %	BV1	5	add.	0.31 (0.33)	0.48 (0.50)	0.69 (0.72)	0.81 (0.85)
32	var	1	1	corr.	5	5 %	2,5 %	BV1	5	add.	0.32 (0.35)	0.45 (0.50)	0.64 (0.71)	0.76 (0.86)
33	var	1	1	corr.	5	5 %	7,5 %	BV1	5	add.	0.27 (0.36)	0.41 (0.51)	0.57 (0.73)	0.62 (0.87)
34	var	1	1	corr.	5	5 %	10 %	BV1	5	add.	0.28 (0.39)	0.41 (0.56)	0.51 (0.74)	0.59 (0.86)
35	var	1	1	corr.	5	5 %	12,5 %	BV1	5	add.	0.22 (0.36)	0.35 (0.55)	0.48 (0.74)	0.60 (0.88)
36	var	1	1	corr.	5	5 %	15 %	BV1	5	add.	0.20 (0.38)	0.31 (0.55)	0.42 (0.75)	0.54 (0.87)
37	var	1	1	corr.	5	5 %	17,5 %	BV1	5	add.	0.19 (0.38)	0.28 (0.57)	0.45 (0.74)	0.56 (0.86)
38	var	1	1	corr.	5	5 %	20 %	BV1	5	add.	0.25 (0.43)	0.33 (0.56)	0.34 (0.75)	0.23 (0.88)
39	var	1	1	corr.	5	5 %	30 %	BV1	5	add.	0.10 (0.42)	0.21 (0.51)	0.42 (0.74)	0.50 (0.85)
40	var	1	1	corr.	5	5 %	40 %	BV1	5	add.	0.20 (0.39)	0.30 (0.54)	0.33 (0.74)	0.27 (0.88)
41	var	1	1	corr.	5	5 %	50 %	BV1	5	add.	0.06 (0.43)	0.24 (0.51)	0.45 (0.65)	0.54 (0.71)
42	var	1	1	corr.	5	5 %	100 %	BV1	5	add.	0.22 (0.35)	0.36 (0.45)	0.49 (0.59)	0.57 (0.70)
43	var	1	1	corr.	5	80 %	5 %	BV1	5	add.	0.27 (0.43)	0.41 (0.55)	0.60 (0.74)	0.69 (0.87)
44	var	2	2	covar.	5	5 %	5 %	BV1	5	mult.	1.70 (1.00)	2.04 (1.00)	2.22 (1.00)	2.25 (1.00)
45	var	1	1	covar.	5	5 %	5 %	BV1	5	mult.	0.68 (1.00)	0.53 (1.00)	0.32 (1.00)	0.37 (1.00)
46	var	2	2	corr.	5	5 %	5 %	BV1	5	mult.	1.59 (1.00)	1.99 (1.00)	2.21 (1.00)	2.25 (1.00)
47	var	1	1	corr.	5	5 %	5 %	BV1	5	mult.	0.69 (1.00)	0.50 (1.00)	0.31 (1.00)	0.36 (1.00)

edge of the Tipping Point in practice). In the case of additive noise, considering previous EOFs like in SD2 may not improve the signal-to-noise ratio because the signal is much weaker away from the Tipping Point. This is of particular importance if the curves of EWS vs. B do not differ substantially between different EOFs. Furthermore, autocorrelations and correlations are more generic indicators while variances can be affected by multiplicative noise in any way. When analysing a system whose variability is of unknown nature, the generic approach is thus probably the most adequate choice. Fig. 10 shows the resulting signal list when using time series of length 100 000 yr and such generic options (referred to as set 1 in Table 4).

Although these options should be applicable to many systems, they may not lead to the most robust results. In system 3 with multiplicative noise, SD2 and ER2 in combination with covariance-based EOFs are of particular advantage. Figure 11 shows the signal list for the multiplicative noise case when using time series of length 10 000 yr and options as set 21 in Table 4. While for the additive noise case the AC's trajectories at the hotspot (red area) and the complete area always look alike (not shown), they differ substantially in the multiplicative case: at the hotspot the signal starts to emerge early, even when projecting on a leading EOF far from the Tipping Point (red curves in Fig. 9). This is not the case for the other areas because the variability of the system differs substantially from the critical mode. As variances at the hotspot are very small, the EOF pronounces other elements than the hotspot and slowing down will thus not be observed in the projection. Close to the Tipping Point, the variance at the hotspot increases not only due to slowing down, but also due to the multiplicative noise which enhances variance as vegetation cover decreases. Therefore, the relative increase in variance is particularly large at the hotspot. Close to the Tipping Point, the system's variability becomes dominated by the critical mode and slowing down can be seen in the complete as well as the hotspot area. By using SD2, ER2 and covariance-based EOFs, we use this property of the system to better distinguish the elements from each other. As a result, the hotspot can be detected much easier than in the additive noise case. Using time series of 10 000 yr each, the hotspot is clearly visible in the signal list for $n_{\max} = 5$ (Fig. 11). Hence, an even more robust hotspot detection can be achieved from time series ten times shorter than in the additive noise case.

In a more general case, additive and multiplicative noise may occur at the same time. In our system 3, the multiplicative noise would dominate the results if noise levels leading to similar variance in V were chosen. However, it is not a priori clear what would happen in other systems whose properties are not well-known. Under such conditions the generic approach using cross- and autocorrelations with SD1 and ER1 would be the safest option in the light of our results.

We now continue with our list of conditions:

- The choice of *time slices* should cover a range of B where the changes in steady state are already pronounced to achieve a good signal-to-noise ratio. In Table 4 we distinguish three different vectors of B-values: BV1: (150, 90, 55, 43), BV2: (300, 200, 100, 75, 43), BV3: (150, 90, 55). For Figs. 7–11 and Table 3 we use BV1 (dashed vertical lines in Fig. 5).
- The *initial threshold*, t_{ini} , and *increment*, t_{inc} , should be chosen small (a few percent of the maximum signal). If they are larger, the calculation is faster and not necessarily worse in performance, but the signal list will not be well sampled. A better sampling of each element's contribution to the signal allows a clearer discrimination between the elements in figures like Figs. 10 and 11. Particularly low η_1 for some larger t_{inc} (Table 4) result from the effect that too many elements are removed at once after increasing the threshold.
- The *maximum number of elements per part*, n_{\max} , can be chosen small for first results. The smaller n_{\max} , the faster the algorithm. When repeating the analysis with larger n_{\max} , the signal list gives an indication of the size of the hotspot (or hotspots). As long as the maximum signal in the list clearly increases with n_{\max} , the number of elements which form a common hotspot is larger than n_{\max} . As Figs. 10 and 11 document, the full hotspot may already be identified for n_{\max} smaller than the hotspot, if t_{ini} and t_{inc} are small to allow a robustly sampled signal list.
- The *length of the time series*, T , as compared to the key variable's timescale τ has a major influence on the method's performance. As the time series provide only a limited sample, the performance will increase with T . If a single available realisation of the time series is too short, the statistical properties of the variations are insufficiently sampled and a hotspot detection can yield wrong results. It should therefore be checked whether the identified hotspot is robust to T by comparing different parts of the time series. Methods of block bootstrapping suited for time series (Politis, 2003) could in principle be applied to the full analysis to derive uncertainty estimates.
- The detectability of a hotspot, given a specific length of the time series, very much depends on intrinsic system properties like its *connectivity and the strength of the destabilising feedback*. The more elements contribute to a hotspot, the more difficult it is to detect. n_{\max} should be chosen large in such a case to determine the large extent of the hotspot which slows down the algorithm. More importantly, the stronger the slowing down and the better the elements can be distinguished, the easier

the hotspot detection. Our system 3 may already provide a rather demanding case as several 10 000–100 000 yr long time slices are required for a robust hotspot detection. This time is rather beyond feasibility for climate models of intermediate or high complexity and a hotspot like in system 3 would hardly be detectable. However, if hotspots of a more pronounced structure exist, they could be detected more easily. As an example, consider the most optimistic case of an univariate process where autocorrelation increases substantially over time. This increase would be detectable within the order of some 100 time steps (Ditlevsen and Johnsen, 2010). Part II of our two-part paper presents a hotspot detection from climate model time series of hundreds to thousands of years length as another example. It is therefore not possible to provide a general statement on the required length of time series. The required length depends on the nature of the potential hotspot, the exact thing that one aims to infer with the analysis. However, this problem does not impose any restrictions to the applicability of the method, but it implies that a negative result can either be due to the non-existence of slowing down at a hotspot or to too short time series.

5 Summary and conclusions

By applying a simple stochastic model, we have demonstrated that EWS at individual elements of a coupled system are no generic precursors of a sudden transition at a tipping point. If the local feedback of a particular element is weak or if the element's tipping is induced by other elements, EWS are not apparent until the bifurcation parameter is very close to its critical point. In this case the signal cannot be called early anymore, and a prediction of a sudden transition, together with the area where it will occur, must fail. On the other hand, we have documented that indicators of slowing down can potentially be used to infer knowledge on the causality of a sudden transition from sufficiently long time series. To this end, we have devised an algorithm to detect the hotspot or hotspots of slowing down in a many-element system. As slowing down indicates a loss in stability of the current state, the detected hotspot indicates a region where the system's susceptibility to perturbations becomes large.

Although our system is meant to represent the vegetation-atmosphere interaction in Northern Africa, the method of analysis is generic in the sense that it can be applied to any system satisfying the basic assumptions common to EWS approaches:

- The system is supposed to be close to a deterministic state (in terms of dynamical systems, a slow manifold), which loses stability.
- The system's variability results from small white noise.

It should be noted that the existence of a bifurcation is not a prerequisite of our method. Even in the case of weaker feedbacks and a more gradual transition will a change in stability be reflected in slowing down. However, the detectability of the signal tends to decrease as compared to a bifurcation where the system approaches a random walk. The main difference to previous applications of EWS is that our method does not only calculate the magnitude of slowing down but also identifies the subsystem where it occurs.

In principle, a prediction of sudden transitions could also be attempted with this approach. As new data points become available, new EOFs and projections may be constructed. As for any prediction based on EWS, it must of course be known in advance which maximum signal is to be expected (Thompson and Sieber, 2011). For example, autocorrelation only comes close to 1 when there is a bifurcation, but peaks at lower values in less extreme cases.

In addition, the very large data requirements imply a vast separation between the timescale of changing external conditions and the intrinsic timescale of the system, a condition that is not often satisfied. Although we focus on autocorrelation and temporal variability, other indicators of slowing down such as spatial variability could be applied within the same iterative framework and may lead to better performances. As our additive and multiplicative noise case illustrate, the more the analysis method is tailored to a specific system, the more a priori knowledge on the data generating process is needed. For example, variance may increase or decrease when approaching a threshold, depending on the system under consideration (Brock and Carpenter, 2010; Dakos et al., 2012).

Additional caveats are imposed by unaccounted or changing properties of the external noise, which would affect EWS (Carpenter and Brock, 2006; Scheffer et al., 2009; Ditlevsen and Johnsen, 2010). In particular, we have only used white noise which is uncorrelated in space. However, it would physically be more reasonable to account for spatial correlations in the atmospheric variability. This could reduce the detectability of hotspots because correlations between the state variables could not be attributed to spatial interactions alone, but would partly result from correlations in the noise.

Other problems may arise in cases of large noise. The local stability of the deterministic state may not be represented well anymore in EWS, and the noise can lead to an early tipping. More fundamentally, the system's mean behaviour in the large noise regime may not reflect its deterministic structure anymore due to noise-induced transitions (Horsthemke and Lefever, 1984). The link between a system's susceptibility and statistical properties of its variability breaks down under such conditions.

Within these limitations, our results suggest an alternative applicability of EWS which may contribute to a better understanding of numerical models. In this regard our study is a concretion of Lenton's recent conclusion: "Even if further research shows that early warning is unachievable in

practice, it could still provide valuable information on the vulnerability of various tipping elements to noise-induced changes.” (Lenton, 2011). To this end, more systematic studies on the performance of indicators of slowing down for different classes of models will be particularly beneficial.

Acknowledgements. S. B. is grateful to Reiner Lauterbach for his very kind and exceptionally clear explanations of mathematical issues, to Jin-Song von Storch for her helpful comments, to Hermann Held for instructive discussions, to Jörn-Peter Boll for creative stimulation and to Walter White for the display of his inspiring career. K. F., Max Planck Fellow, acknowledges the support by the Max Planck Society. This work is funded by Deutsche Forschungsgesellschaft, Cluster of Excellence “CliSAP” (DFG EXC 177).

Edited by: N. de Noblet

The service charges for this open access publication have been covered by the Max Planck Society.

References

- Bathiany, S., Claussen, M., and Fraedrich, K.: Implications of climate variability for the detection of multiple equilibria and for rapid transitions in the atmosphere-vegetation system, *Clim. Dynam.*, 38, 1775–1790, doi:10.1007/s00382-011-1037-x, 2012.
- Bathiany, S., Claussen, M., and Fraedrich, K.: Detecting hotspots of atmosphere-vegetation interaction via slowing down – Part 2: Application to a global climate model, *Earth Syst. Dynam.*, 4, 79–93, doi:10.5194/esd-4-79-2013, 2013.
- Brock, W. A. and Carpenter, S. R.: Interacting regime shifts in ecosystems: implication for early warnings, *Ecol. Monogr.*, 80, 353–367, doi:10.1890/09-1824.1, 2010.
- Brovkin, V., Claussen, M., Petoukhov, V., and Ganopolski, A.: On the stability of the atmosphere-vegetation system in the Sahara/Sahel region, *J. Geophys. Res.-Atmos.*, 103, 31613–31624, doi:10.1029/1998JD200006, 1998.
- Brovkin, V., Bendtsen, J., Claussen, M., Ganopolski, A., Kubatzki, C., Petoukhov, V., and Andreev, A.: Carbon cycle, vegetation, and climate dynamics in the Holocene: experiments with the CLIMBER-2 Model, *Global Biogeochem. Cy.*, 16, 1139, doi:10.1029/2001GB001662, 2002.
- Carpenter, S. R. and Brock, W. A.: Rising variance: a leading indicator of ecological transition, *Ecol. Lett.*, 9, 308–315, doi:10.1111/j.1461-0248.2005.00877.x, 2006.
- Charney, J. G.: Dynamics of deserts and drought in the Sahel, *Q. J. Roy. Meteor. Soc.*, 101, 193–202, doi:10.1002/qj.49710142802, 1975.
- Claussen, M.: Modeling bio-geophysical feedback in the African and Indian monsoon region, *Clim. Dynam.*, 13, 247–257, doi:10.1007/s003820050164, 1997.
- Claussen, M.: On multiple solutions of the atmosphere-vegetation system in present-day climate, *Global Change Biol.*, 4, 549–559, doi:10.1046/j.1365-2486.1998.t01-1-00122.x, 1998.
- Claussen, M.: Late Quaternary vegetation-climate feedbacks, *Clim. Past*, 5, 203–216, doi:10.5194/cp-5-203-2009, 2009.
- Claussen, M., Kubatzki, C., Brovkin, V., Ganopolski, A., Hoelzmann, P., and Pachur, H. J.: Simulation of an abrupt change in Saharan vegetation in the mid-Holocene, *Geophys. Res. Lett.*, 26, 2037–2040, doi:10.1029/1999GL900494, 1999.
- Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V., and Held, H.: Slowing down as an early warning signal for abrupt climate change, *Proc. Natl. Acad. Sci. USA*, 105, 14308–14312, doi:10.1073/pnas.0802430105, 2008.
- Dakos, V., van Nes, E. H., Donangelo, R., Fort, H., and Scheffer, M.: Spatial correlation as leading indicator of catastrophic shifts, *Theor. Ecol.*, 3, 163–174, doi:10.1007/s12080-009-0060-6, 2010.
- Dakos, V., Kefi, S., Rietkerk, M., van Nes, E. H., and Scheffer, M.: Slowing down in spatially patterned ecosystems at the brink of collapse, *Am. Nat.*, 177, E153–E166, doi:10.1086/659945, 2011.
- Dakos, V., van Nes, E. H., D’Odorico, P., and Scheffer, M.: Robustness of variance and autocorrelation as indicators of critical slowing down, *Ecology*, 93, 264–271, doi:10.1890/11-0889.1, 2012.
- Ditlevsen, P. D. and Johnsen, S. J.: Tipping points: early warning and wishful thinking, *Geophys. Res. Lett.*, 37, L19703, doi:10.1029/2010GL044486, 2010.
- Donangelo, R., Fort, H., Dakos, V., Scheffer, M., and Van Nes, E. H.: Early warnings for catastrophic shifts in ecosystems: comparison between spatial and temporal indicators, *Int. J. Bifurcat. Chaos*, 20, 315–321, doi:10.1142/S0218127410025764, 2010.
- Efron, B.: 1977 Rietz Lecture – Bootstrap Methods – Another Look At the Jackknife, *Ann. Stat.*, 7, 1–26, doi:10.1214/aos/1176344552, 1979.
- Guttal, V. and Jayaprakash, C.: Changing skewness: an early warning signal of regime shifts in ecosystems, *Ecol. Lett.*, 11, 450–460, doi:10.1111/j.1461-0248.2008.01160.x, 2008.
- Guttal, V. and Jayaprakash, C.: Spatial variance and spatial skewness: leading indicators of regime shifts in spatial ecological systems, *Theor. Ecol.*, 2, 3–12, doi:10.1007/s12080-008-0033-1, 2009.
- Hales, K., Neelin, J. D., and Zeng, N.: Sensitivity of tropical land climate to leaf area index: role of surface conductance versus albedo, *J. Climate*, 17, 1459–1473, doi:10.1175/1520-0442(2004)017<1459:SOTLCT>2.0.CO;2, 2004.
- Hasselmann, K.: Stochastic climate models. Part I. Theory, *Tellus*, 28, 473–485, doi:10.1111/j.2153-3490.1976.tb00696.x, 1976.
- Held, H. and Kleinen, T.: Detection of climate system bifurcations by degenerate fingerprinting, *Geophys. Res. Lett.*, 31, L23207, doi:10.1029/2004GL020972, 2004.
- Horsthemke, W. and Lefever, R.: *Noise-Induced Transitions*, Springer, 1984.
- Irizarry-Ortiz, M. M., Wang, G. L., and Eltahir, E. A. B.: Role of the biosphere in the mid-Holocene climate of West Africa, *J. Geophys. Res.-Atmos.*, 108, 4042, doi:10.1029/2001JD000989, 2003.
- Kleinen, T., Held, H., and Petschel-Held, G.: The potential role of spectral properties in detecting thresholds in the Earth system: application to the thermohaline circulation, *Ocean Dynam.*, 53, 53–63, doi:10.1007/s10236-002-0023-6, 2003.
- Lenton, T. M.: Early warning of climate tipping points, *Nature Clim. Change*, 1, 201–209, doi:10.1038/NCLIMATE1143, 2011.
- Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S., and Schellnhuber, H. J.: Tipping elements in the Earth’s climate system, *Proc. Natl. Acad. Sci. USA*, 105, 1786–1793,

- doi:10.1073/pnas.0705414105, 2008.
- Lenton, T. M., Livina, V. N., Dakos, V., and Scheffer, M.: Climate bifurcation during the last deglaciation?, *Clim. Past*, 8, 1127–1139, doi:10.5194/cp-8-1127-2012, 2012a.
- Lenton, T. M., Livina, V. N., Dakos, V., van Nes, E. H., and Scheffer, M.: Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness, *Phil. Trans. R. Soc. A*, 370, 1185–1204, doi:10.1098/rsta.2011.0304, 2012b.
- Liu, Z. Y., Wang, Y., Gallimore, R., Notaro, M., and Prentice, I. C.: On the cause of abrupt vegetation collapse in North Africa during the Holocene: climate variability vs. vegetation feedback, *Geophys. Res. Lett.*, 33, L22709, doi:10.1029/2006GL028062, 2006.
- Politis, D. N.: The impact of bootstrap methods on time series analysis, *Stat. Sci.*, 18, 219–230, doi:10.1214/ss/1063994977, 2003.
- Priestley, M. B.: *Spectral Analysis and Time Series. Volume 1: Univariate Series*, Academic Press, 1981.
- Rietkerk, M., Stroosnijder, L., Prins, H., and Ketner, P.: Sahelian rangeland development: a catastrophe?, *J. Range Manage.*, 49, 512–519, doi:10.2307/4002292, 1996.
- Scheffer, M., Bascompte, J., Brock, W. A., Brovkin, V., Carpenter, S. R., Dakos, V., Held, H., van Nes, E. H., Rietkerk, M., and Sugihara, G.: Early-warning signals for critical transitions, *Nature*, 461, 53–59, doi:10.1038/nature08227, 2009.
- Sieber, J. and Thompson, J. M. T.: Nonlinear softening as a predictive precursor to climate tipping, *Phil. Trans. R. Soc. A*, 370, 1205–1227, doi:10.1098/rsta.2011.0372, 2012.
- Strogatz, S.: *Nonlinear Dynamics and Chaos. With Applications to Physics, Biology, Chemistry, and Engineering*, Perseus Books Group, 1994.
- Thompson, J. M. T. and Sieber, J.: Climate tipping as a noisy bifurcation: a predictive technique, *IMA J. Appl. Math.*, 76, 27–46, doi:10.1093/imamat/hxq060, 2011.
- Wang, G. L.: A conceptual modeling study on biosphere-atmosphere interactions and its implications for physically based climate modeling, *J. Climate*, 17, 2572–2583, doi:10.1175/1520-0442(2004)017<2572:ACMSOB>2.0.CO;2, 2004.
- Wang, G. L. and Eltahir, E. A. B.: Biosphere-atmosphere interactions over West Africa. II: Multiple climate equilibria, *Q. J. Roy. Meteor. Soc.*, 126, 1261–1280, doi:10.1002/qj.49712656504, 2000.
- Wiesenfeld, K.: Noisy precursors of nonlinear instabilities, *J. Stat. Phys.*, 38, 1071–1097, doi:10.1007/BF01010430, 1985a.
- Wiesenfeld, K.: Virtual Hopf phenomenon – a new precursor of period-doubling bifurcations, *Phys. Rev. A*, 32, 1744–1751, doi:10.1103/PhysRevA.32.1744, 1985b.
- Wiesenfeld, K. and McNamara, B.: Small-signal amplification in bifurcating dynamical systems, *Phys. Rev. A*, 33, 629–642, doi:10.1103/PhysRevA.33.629, 1986.
- Wissel, C.: A universal law of the characteristic return time near thresholds, *Oecologia*, 65, 101–107, doi:10.1007/BF00384470, 1984.
- Zeng, N. and Neelin, J. D.: The role of vegetation-climate interaction and interannual variability in shaping the African savanna, *J. Climate*, 13, 2665–2670, doi:10.1175/1520-0442(2000)013<2665:TROVCI>2.0.CO;2, 2000.