National Physical Laboratory

NPL Report MS 47

TIPPING POINT ANALYSIS AND ITS APPLICATIONS

VALERIE LIVINA

MARCH 2023

Department for
Science, Innovation
& Technology

NPL Report MS 47

Tipping Point Analysis and Its Applications

Valerie Livina

Data Analytics & Modelling Group
Data Science Department

ABSTRACT

The report summaries tipping point analysis techniques for time series of dynamical systems and provides case studies in various applications. The work on tipping point analysis started in 2006 at the University of East Anglia (funded by Natural Environment Research Council and AXA Research Fund) and has been conducted since 2012 at National Physical Laboratory (funded by National Measurement System).

ISSN 1754-2960

https://doi.org/10.47120/npl.MS47

National Physical Laboratory
Hampton Road, Teddington, Middlesex, TW11 0LW

This work was funded by the UK Government's Department for Science, Innovation & Technology through the UK's National Measurement System programmes.

Approved on behalf of NPLML by Peter Harris,

Science Area Leader

Contents

List of Acronyms

ACF	Autocorrelation Function
ACF(1)	Lag-1 Autocorrelation
AHU	Air Handling Unit
AR(p)	Autoregressive Model (of order p)
ARIMA	Autoregressive Integrated Moving Average
CTBTO	Comprehensive Nuclear-Test-Ban Treaty Organization
CTE	Coefficient of Thermal Expansion
DFA	Detrended Fluctuation Analysis
DMSP	Defense Meteorological Satellite Program
DO	Dansgaard-Oeschger event
ENSO	El Niño/Southern Oscillation
EOF	Empirical Orthogonal Function
EWS	Early Warning Signal
FFT	Fast Fourier Transform
GICC05	Greenland Ice-Core Chronology 2005
GRIP	Greenland Ice Core Project
IPCC	The Intergovernmental Panel on Climate Change
LGM	Last Glacial Maximum
LSTM	Long Short-Term Memory
NASA	The National Aeronautics and Space Administration
NGRIP	North Greenland Ice Core Project
NOAA	National Oceanic and Atmospheric Administration
NS	Nash-Sutcliffe efficiency
OAT	Outside Air Temperature
ONI	Oceanic Niño Index
PB	Percent Bias
PCA	Principal Component Analysis
PDF	Probability Density Function
PS	Power Spectrum
RMSE	Root Mean Squared Error
SAT	Supply Air Temperature
SMMR	Scanning Multichannel Microwave Radiometer
SOI	Southern Oscillation Index
SPL	Sound Pressure Level
SSA	Singular Spectrum Analysis
SSM/I	Special Sensor Microwave/Imager
SSMIS	Special Sensor Microwave Imager/Sounder
UKF	Unscented Kalman Filter

Chapter 1

Introduction

The origin of tipping points of dynamical systems can be traced to works of Henri Poincare who developed the bifurcation theory [Poincare, 1892]. Tipping points of the Earth system are currently an important topic of research in climatology, with increasing number of applications of mathematical bifurcation theory.

Nowadays, climatologists are well familiar with the Poincaré-Andronov-Hopf bifurcation observed in climatic variables. In the 1960-1970s, an intuitive idea of a bifurcation appeared in social sciences as the term ”tipping point” mentioned by many sociologists, such as Everett Rogers, Morton Grodzins, Thomas Schelling, and others. Understanding of tipping in the context of social sciences was within a general idea of small perturbations of a balanced system that may lead to its toppling. That time, no analytical sociophysics models existed to link the idea of social tipping with mathematical works of Poincare.

Malcolm Gladwell went further with developing this trend in his best-selling book ”The Tipping Point: How Little Things Can Make a Big Difference” [Gladwell, 2000]: he expanded the approach to biophysical systems (for example, to spread of sexually transmitted diseases), but he still used the language of social science to formulate ”the three rules of the tipping point: the Law of the Few, the Stickiness Factor, the Power of Context”.

But it was not long before the idea of tipping was brought back to natural sciences: [Lenton et al., 2008] published the seminal paper on tipping elements in the Earth system, in which the main tipping elements — vulnerable regions and subsystems of the climate system — were identified and described. This work initiated multiple publications in geophysics and paleoclimate, focussing on early warning signals of tipping points. They were developing mathematical techniques that would allow one to detect and quantify upcoming transitions and/or bifurcations in a time series of a climatic variable. Such critical transitions are signatures of tipping element, i.e., climate subsystems and regions experiencing tipping.

Mathematically, tipping points are critical transitions and bifurcations in time series [Livina et al., 2011] that can be analysed using the techniques described in this report. Tipping point analysis can be applied to a broad range of dynamical systems whose trajectories (time series) experience critical behaviour, and in this report several examples of such systems with tipping points will be analysed.

A critical transition may be a drift or a change point (in mean value or amplitude) without structural changes of the system potential. The bifurcation, however, leads to a change of the structure of the potential, with appearing or disappearing potential wells [Livina et al., 2010]. This is illustrated in Figure. 1.1.

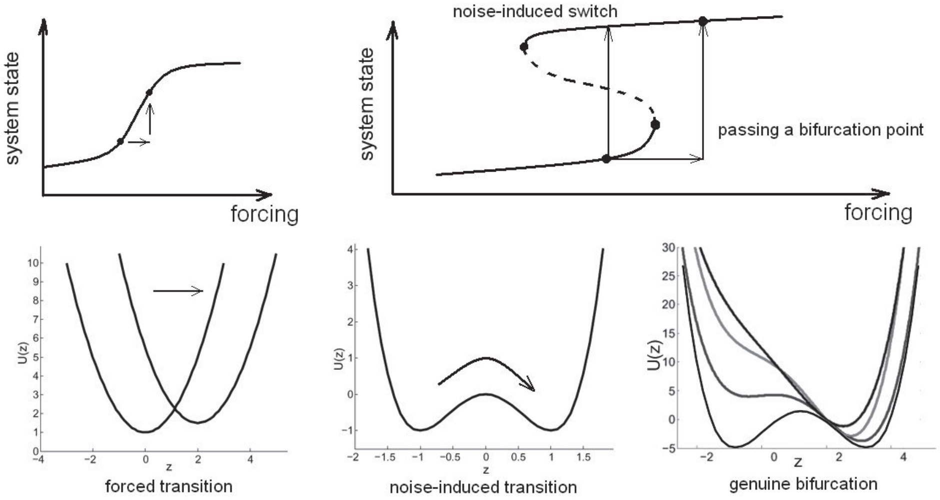

Figure 1.1: Schematic of transitions and bifurcations: upper panels — bifurcation diagrams describing transitions and bifurcations; bottom panels — potential curves.

[Livina et al., 2011] identifies three types of system changes which may be the cause of a tipping event:

1. A forced transition, where the potential function stays the same shape, but is moved, so that there is a trend (or possibly a sudden jump) in the time series.
2. A noise-induced transition, where the system shifts from one stable state to another due to a large initial perturbation.
3. A genuine bifurcation, where the shape of the potential function is changed.

This framework is based on statistical physics (governed by the Fokker-Planck equation) and can be applied to a wide range of dynamical systems with recorded trajectories. In the context of mechanical or electromagnetic systems, such transitions and bifurcations typically denote system failures. Such failures in such systems can be analysed using the techniques of tipping point analysis, which helps anticipate them.

Analysis of early warning signals of tipping points has generated a large number of impactful research papers, for instance, in ecology (see [Scheffer et al., 2009] and references therein), although the earlier studies of noisy bifurcation precursors appeared already in the 1980s [Wiesenfeld, 1985]. Although well-developed, these classical theories are difficult to apply in data-driven analysis of dynamical systems with unknown analytical description. Moreover, societally-relevant transitions (without structural change of the dynamical system), may be as important as bifurcations and require adequate analytical tools and techniques of analysis.

Although there is common understanding of tipping points as critical changes detectable in time series, and attempts are made to classify tipping points [Ashwin et al., 2012, Thompson and Sieber, 2011], there is a gap between the contemporary research literature on early warning signals that are being analysed with basic autocorrelation techniques and analytical bifurcation theory in dynamical systems of the second half of the 20th century [Arnold, 1986, Milnor, 1963]. Tipping point analysis fills this gap and provides flexible instruments (analytical tools as well as software) for such analysis. The approach combines a general analytical description of the dynamical system with data-driven derivation of the underlying system potential, which makes it flexible and applicable to a broad range of real-world dynamical systems.

One of the main reasons for that is that many models employed in the theory of bifurcations are purely deterministic (although chaotic), which makes them of restricted applicability for studying various real-world complex systems, where there is lack of knowledge of analytical description of the system. For example, climate system modelling, which started as physically-based and deterministic, with years of development involved a higher and higher number of variables and parameters, and at a certain level of complexity it was understood that its modelling performance can be matched by simpler and computationally lighter stochastic models [Hasselmann, 1976, Palmer and Williams, 2008]. In fact, this holds not only for climatic models, but also for models of many other complex systems.

A stochastic model involving deterministic and stochastic components (static or dynamic) is a powerful yet simple tool for modelling arbitrary time series of real-world dynamical systems. These systems are characterised by complexity and a high number of variables producing a single system trajectory (for example, a temperature record influenced by hundreds of other climate variables like humidity, wind speed, etc). In most cases, it is impossible to reconstruct a deterministic model from such recorded datasets, and the analysis of the system has to be data-driven. This necessaity makes analysis of tipping points even more complicated, and novel techniques are needed to address this challenge. The main advantage of such stochastic approaches is that the multitude of the complex factors influencing the dynamics of the system are separated into two major components, with smaller-scale random fluctuations under the umbrella of the stochastic term. This allows one to monitor the major variability of the system state, combined with analysis of the shorter-term fluctuation dynamics.

Given one-dimensional trajectory of a dynamical system (recorded time series), the system dynamics may be modelled by the following, very general, stochastic equation with state variable z and time t:

$$\dot{z} = D(z,t) + S(z,t), \tag{1.1}$$

where D and S are deterministic and stochastic components, correspondently. Component $D(z,t)$ may be stationary or include a trend, depending on the studied system. The probability distribution of such a system, however complex it may be, in the majority of cases can be approximated by a polynomial of even order (so-called potential). The stochastic component, in the simplest case, may be Gaussian white noise, although in real systems it is often more complex than that, possessing correlations (red noise), anti-correlations (blue noise), multifractal and other nonlinear properties.

Tipping point analysis consists of the following three stages:

- anticipating (pre-tipping)
- detecting (tipping)

- forecasting (post-tipping)

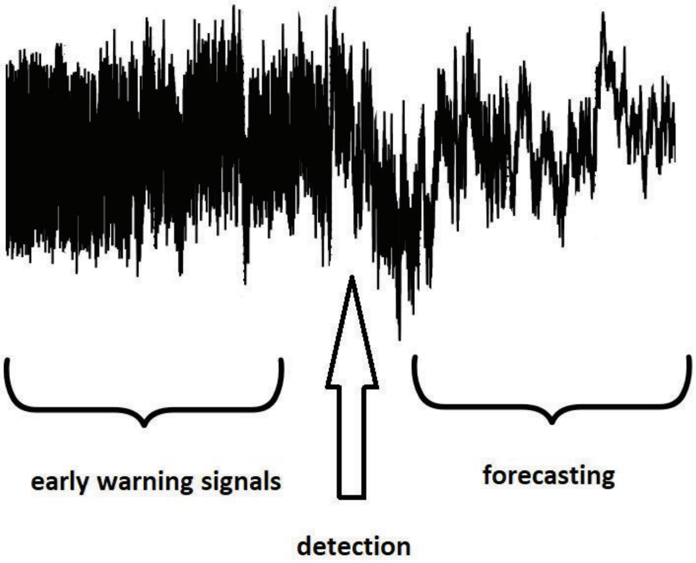

Figure 1.2: Schematic of a time series experiencing a tipping point with change from stationary to non-stationary behaviour, and different stages of tipping point analysis where various machine learning techniques can be applied.

The schematic in Figure 1.2 illustrates how various techniques can be used at different stages of a failure in a single time series. The time series is stationary in the beginning, then experiences a tipping point in the middle and later oscillates in a non-stationary way.

Tipping points can be defined and classified in terms of the underlying system potential whose derivative, if it exists, defines the deterministic term in Equation (1.1). If the potential structure (number of potential wells) changes, the tipping point is a genuine bifurcation. If the potential structure remains the same, while the trajectory of the system samples various states, the tipping point is transitional. An example of such a transition is the record of global temperature, which has the same structure of fluctuations with drifting (under forcing or noise-induced). It is important to identify and distinguish types of tipping points (bifurcations vs transitions) for the purpose of the further forecast. This is a key question in order to define the horizon of the forecast and general predictability of the tipping point (whether it is abrupt or gradual).

Single-variable datasets (time series, or trajectories of dynamical systems) were studied in this context, applying early warning (EWS) signal techniques in geophysics [Livina and Lenton, 2007], structural health monitoring [Livina et al., 2014, Perry et al., 2016] and in studying semiconductor devices [Livina et al., 2020]. However, many dynamical systems nowadays provide measurements of multiple variables and multiple modes of state observation. This poses the challenge of analysing multidimensional datasets to identify if a critical behaviour starts appearing in the system, which may lead to an eventual failure. One approach is to apply techniques of dimension reduction, such as principal component analysis [Prettyman et al., 2018]. However, in some systems various measurements may not be simultaneous, and further development of multivariate analysis is necessary. In particular, clustering time series often helps to identify groups of records with similar properties, based on the chosen distance measure.

When applying such data-driven techniques, there are several important things to keep in mind. First, a forecast model would be more accurate if longer training data is provided. Second, the accuracy of any forecast depends on the forecast horizon, i.e., how far into future one can attempt to run the model. Third, the dynamics of changes — abrupt or gradual — would affect detection and forecast of time series.

As the system dynamics unfolds, and the data is investigated for possible tipping, one can attempt to identify early warning signals of a critical transition. Furthermore, failure can be detected at the moment of its happening, using several techniques of time series analysis, such as Canonical Correlation Analysis and potential analysis. Finally, when the failure is certain, one can attempt to forecast its evolution when possible; in some cases, however, it may be not possible due to a quickly evolving system or an inadequate sampling rate.

The techniques used for early warning signals, detection and forecasting of tipping events, are illustrated in the flowchart in Figure 1.3.

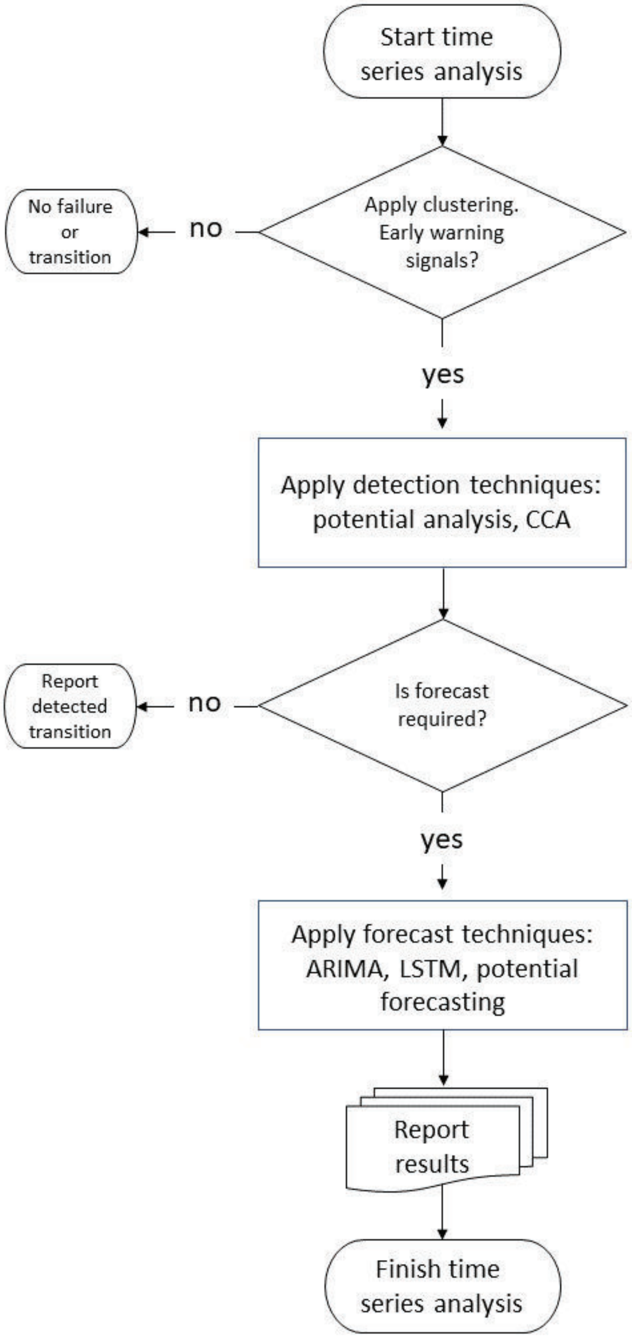

Figure 1.3: Flowchart of the tipping point analysis techniques.

Given a set of time series data (e.g. a measurement variable of a physical system or computational model) where a transition or bifurcation has occurred, one may detect an Early Warning Signal, which should be measurable before the tipping point. This EWS can then be used to predict future tipping points.

[Dakos et al., 2008] lists several EWS techniques which have been attempted: analysis of spectral properties [Kleinen et al., 2003], degenerate fingerprinting [Held and Kleinen, 2004] and a modification of Detrended Fluctuation Analysis (DFA) [Livina and Lenton, 2007].

A key concept involved in these methods is that of resilience or recovery rates [Scheffer et al., 2001, Veraart et al., 2012], i.e., how quickly the system returns to the stable state having received a small perturbation. Theory states that the recovery rate will slow down as the system approaches a bifurcation point [Scheffer et al., 2009], known as critical slowing down.

The hypothesis is that this critical slowing down will be seen in a wide variety of dynamical systems approaching a bifurcation.

[Van Nes and Scheffer, 2007] test this hypothesis with a number of ecological models and find that the intensity of the critical slowing down effect is linearly, or almost linearly, related to the distance from the tipping point in all cases.

As an example, the following system can be considered:

$$\dot{z}(t) = -\frac{\partial}{\partial z}\left(z^4 + (3-t)\,z^2\right) + \eta_t \tag{1.2}$$

where η_t are independent and Gaussian. This system oscillates within a single well of attraction (at $z = 0$) for $t \leq 3$, at $t = 3$ a bifurcation occurs creating a double well potential system for $t > 3$. The shape of the potential at $t = 0$, $t = 2$ and $t = 3$ is shown in Figure 1.4. Intuitively, the system will take longer to return to the $z = 0$ equilibrium, following a small perturbation, as the sides of the potential well become less steep.

This system is generic and describes a wide range of complex dynamical systems observed in nature. It is an example of the Langevin equation and provides a testbed for several techniques of tipping point analysis, as will be shown in the following sections.

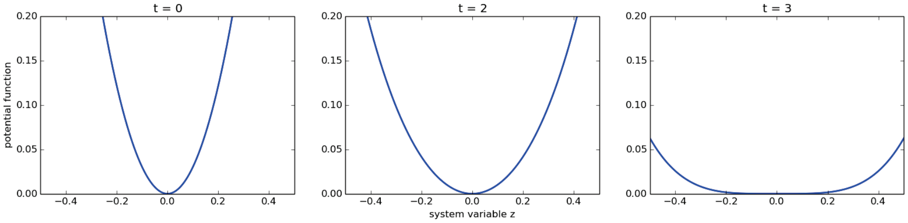

Figure 1.4: The shape of the system potential of the system in equation (1.2) at $t = 0, 2, 3$. The sides of the potential well become less steep as the system approaches the bifurcation.

Chapter 2

Univariate tipping points

2.1 Anticipating Tipping: Early Warning Signals

Temporal scaling properties of time series (or self-similarity, or persistence) can be measured using three techniques: the autocorrelation function (ACF), Detrended Fluctuation Analysis and the power spectrum [Bak et al., 1988, Kantelhardt et al., 2001]. The scaling exponents of these methods can be used for calculation of indicators of Early Warning Signals in time series of dynamical systems.

Real dynamical systems carry memory (e.g., air temperature, SST, river flux), caused by different types of inertia or trends. Physically, this temporal scaling means that the state of the system depends not only on the previous time step, but on the history with some time lag. Statistically, it is described in terms of correlations, and one of the first studies of correlations in climate can be found in [Koscielny-Bunde et al., 1998]. A time series that possesses long-range temporal correlations can be characterized by scaling laws with a relevant scaling exponent, and there exist several methods for estimating the correlations.

The power spectrum $S(f)$, where f is a frequency, in a long-range correlated time series has a power law form with a scaling exponent β (i.e., the power spectrum is proportional to $1/f^{\beta}$). Similarly, the auto-correlation function $C(s)$, where s is a time lag, provides power-law exponent γ. Trends that exist in climate records may cause inaccuracies in measuring these two exponents. Hence more advanced techniques that exclude trends from the data have been developed. One of the methods that copes well with non-stationarity is the Detrended Fluctuation Analysis (DFA).

2.1.1 Autocorrelation exponent

The autocorrelation function (ACF) of a time series, with lag l, is estimated as

$$\text{ACF}_l(X) = \frac{1}{(N-l)s^2} \sum_{j=1}^{N-l} (X_j - \bar{X})(X_{j+l} - \bar{X}) \tag{2.1}$$

where $\bar{X}$ and s are the mean and standard deviation of the series. The ACF scaling exponent γ is defined as the power-law decay rate of the autocorrelation function with increasing lag [Kantelhardt et al., 2001]:

$$ACF_l \sim l^{-\gamma}, \tag{2.2}$$

for long-range correlations. It is possible to estimate the exponent γ by calculating the ACF with a range of lags $l = 1, 2, 3, \ldots,$ and plotting the results on logarithmic axes. The exponent γ can be found by calculating the negative slope of the linear fit to the log-log ACF plot.

Lag-1 ACF, or ACF(1), is used in early warning signal contexts rather than the ACF scaling exponent. Where the time series is long-range correlated, the ACF(1) may be used as a proxy for the ACF scaling exponent: Figure 2.4 shows the ACF exponents and the ACF(1) calculated for different noise series. Where the time series is not long-range correlated, ACF_l decays exponentially with increasing lag and only ACF(1) is indicative of data variability in proximity to a tipping point. Other measures of autocorrelation, such as the Mann-Kendall coefficient, may be used as well [Yue et al., 2002].

2.1.2 Detrended fluctuation analysis exponent

Detrended Fluctuation Analysis, defined in [Kantelhardt et al., 2001], is used by [Livina and Lenton, 2007] as an EWS, similarly to autocorrelation. In this method, the time series x_i is integrated, and the obtained profile (i.e., discrete integral/cumulative sum of x_i) is divided into windows of size s. Next, within each window, the best polynomial fit of chosen order is calculated, and the square of the difference between the polynomial and profile function is evaluated. At the next step, the obtained values are averaged over all windows and the procedure is repeated for different window scales s, arriving at fluctuatons:

$$F^2(\nu, s) = \frac{1}{s} \sum_{i=1}^{s} \left[Y((\nu - 1)s + i) - y_\nu(i) \right]^2, \tag{2.3}$$

where Y is the profile function of initial time series x_i, $y_\nu(i)$ is the best fitting polynomial of order k in segment ν, where ν is the index of non-overlapping windows of size s, and N is the length of the time series. Then over all segments ν is averaged, with a square root to obtain the fluctuation function $F(s)$.

If the series is long-term power-law correlated, the DFA fluctuation function increases by a power-law:

$$F(s) \propto s^{\alpha}, \tag{2.4}$$

where α is the DFA scaling exponent. For uncorrelated records, $\alpha = 0.5$, while for long-term correlated (persistent) records, $\alpha > 0.5$ Kantelhardt et al. [2001]. In DFA of k^{th} order, trends of polynomial order $k - 1$ are eliminated. The DFA exponent α is related to the power spectrum exponent β from $S(f) \sim 1/f^{\beta}$ by $\beta = 2\alpha - 1$, such that for white noise with $\alpha = 0.5$, the power spectrum is flat with $\beta = 0$, and for a random walk signal (integrated white noise) with exponent $\alpha = 1.5$, the power spectrum indicates red noise with $\beta = 2$. The auto-correlation exponent γ in $C(s) \propto s^{-\gamma}$ (here s is the ACF lag) is related to the DFA exponent α by $\alpha = 1 - \gamma/2$, where $0 < \gamma < 1$ (stationary series).

By estimating the slope of the log-log plot of DFA using a linear fit, one calculates the exponent α. Order-2 polynomials are regularly used for the detrending step of DFA for the purpose of EWS detection. Following the approach of [Livina and Lenton, 2007], the DFA exponent is estimated in the temporal range $10 \leq s \leq 100$.

The advantage of DFA is that it avoids artifacts (short-term 'weather noise'), because the DFA exponent is measured in time scale above 10 time units. In contrast, the conventional ACF lag-1 correlations are related to the shortest time scale and are more vulnerable to the influence of noise, and therefore the aggregation of the series is required. Therefore, the series should be long enough to allow aggregation, and short records cannot be studied with the ACF-based technique.

When compared to the auto-correlation power-law exponent γ, the DFA does not have boundary restrictions (imposed on γ by the stationarity condition) and it therefore provides more flexibility in the estimation of correlations. The series should have at least several hundred points per window, and a sufficient number of sliding windows to properly estimate the evolution of the indicator. Therefore, the record should be at least of order a thousand points. To detect an incipient bifurcation, one could monitor the dynamics of the DFA exponent α towards its critical value 1.5 which is approximately equivalent to critical c value 1 (non-stationary state, which leads to bifurcation).

2.1.3 Power spectrum exponent

The power spectrum scaling exponent β is calculated by estimating the slope of the power spectrum $S(f)$ of the data, plotted on logarithmic axes [Bak et al., 1988], from which one can obtain β via the scaling relationship

$$S(f) \sim f^{-\beta}. \tag{2.5}$$

The frequency range $10^{-2} \leq f \leq 10^{-1}$ corresponds to the time range of 10 to 100 units, to put the technique in agreement with with the ACF and DFA methods, as above.

There exist various methods for the calculation of the power spectrum. One can use the periodogram, obtained from the absolute value of the fast Fourier transform.

The PS-exponent can be calculated by log-log plotting the standard periodogram of the data, simply the coefficients given by the fast Fourier transform (FFT). Another method of power spectrum estimation is Welch's method [Welch, 1967]. A line is fitted to the power spectrum estimate (in our case the periodogram) in the frequency range 10^{-2} to 10^{-1}, consistent with the range of 10 segments to 100 segments over which the DFA coefficient is measured [Livina and Lenton, 2007]. The negative of the gradient of the line is the PS exponent.

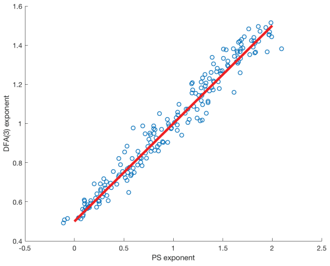

Figure 2.1: long range

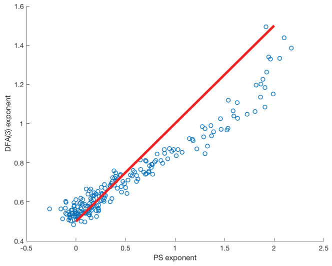

Figure 2.2: short range

[Livina et al., 2012] compared the performance ACF and DFA degenerate fingerprinting methods on artificial data. The ACF and DFA methods are tested on an artificial time series produced by concatenating several series of increasingly correlated noise, from pure white noise to Brownian motion.

2.1.4 Relationships between scaling exponents

These three methods are illustrated in Figure 2.3, where a time series of artificial data and three scaling curves are plotted. The data is a length 10000 red-noise series, with power spectrum scaling exponent of 0.85.

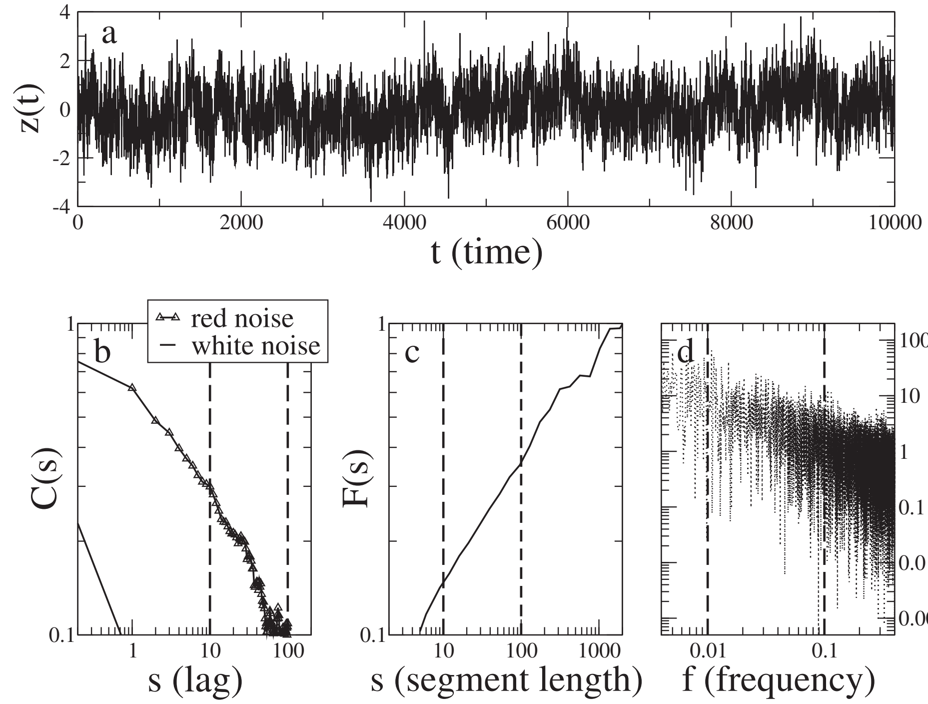

Figure 2.3: Analysis of artificial red noise with scaling exponents calculated using three different methods. (a) Red noise with a power spectrum scaling exponent of 0.85. (b) The ACF of the red noise data is calculated for different lags and the exponent (negative slope) measured in the range $10 \leq s \leq 100$ (dashed lines). Note that the ACF(1) indicator ($C(1)$) is 0.64. The ACF of a white noise series is also plotted for comparison, in this case $C(s) = 0$ for $s \geq 1$ and the exponent is also zero. (c) DFA calculated for the data and the exponent (slope) measured in the range $10 \leq s \leq 100$. (d) The power spectrum of the data, and the exponent (negative slope) measured in the frequency range $10^{-2} \leq f \leq 10^{-1}$.

Analytically, the three scaling exponents have the linear relationship:

$$\alpha = \frac{1+\beta}{2} = 1 - \frac{\gamma}{2} \tag{2.6}$$

in the asymptotic range [Heneghan and McDarby, 2000, Kantelhardt et al., 2001]. A measurement of β could be used in an application where α or γ are conventionally used, to provide an alternative metric for EWS, which may be informative in cases where other indicators fail to detect an approaching tipping point.

2.1.5 Use of scaling exponents as early warning signals

The use of autocorrelation as an EWS is justified by modelling a dynamical system using a one-dimensional autoregressive system AR(1):

$$y_{n+1} = e^{\kappa \Delta t} y_n + \eta_n \tag{2.7}$$

[Held and Kleinen, 2004, Scheffer et al., 2009]. Where η_n are independent and Gaussian, the state $y_n = 0$ is the equilibrium. The system returns to equilibrium exponentially with rate κ (the decay rate). If critical slowing down is simulated by decreasing κ, the autocorrelation coefficient $c \equiv e^{\kappa \Delta t}$ increases, $c \xrightarrow[\kappa \to 0]{} 1$. It is also shown that as the autocorrelation increases so does the variance, thus detecting an increase in variance provides another EWS.

Note that for the purposes of EWS analysis, it is possible to follow the conventional general procedure of fitting an auto-regressive model to the dataset under consideration, but the lag-1 autocorrelation value can also be derived directly from the time series. For AR(1) process $X_t = \phi_1 X_{t-1} + \epsilon_t$, autocovariance $\gamma_k = cov(X_t, X_{t-k}) = cov(\phi_1 X_{t-1} + \epsilon_t, X_{t-k})) = \phi_1 \gamma_{k-1}$. This means that $\gamma_k = \phi_1^{|k|} \gamma_0,\ k \in N$. Since the empirical autocorrelation function is always normalised (i.e., $\gamma_0 = 1$), it means that lag-1 autocorrelation γ_1 coincides with the AR(1) parameter ϕ_1 and can be obtained directly from the ACF of the time series.

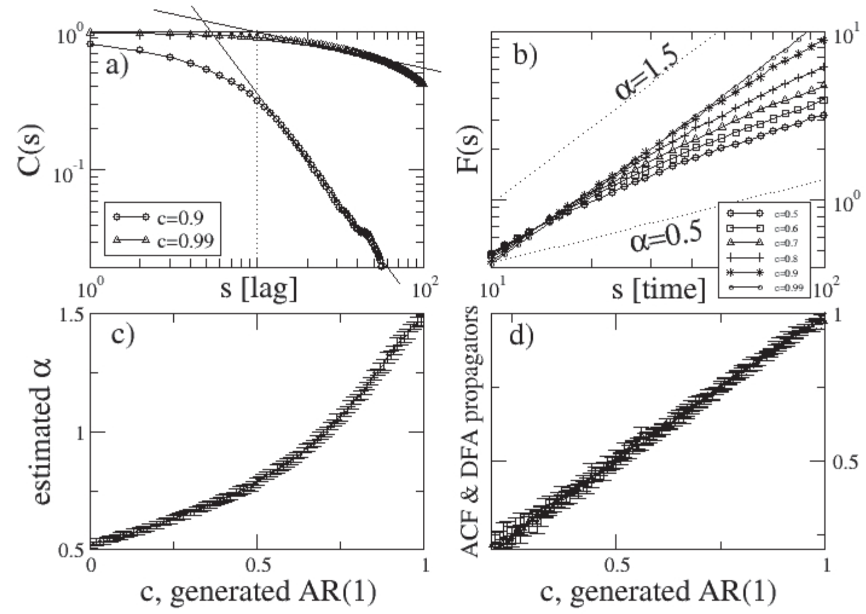

Figure 2.4: a) Auto-correlation function for AR(1) data, $c = 0.9, 0.99$; as $c \to 1$, the exponential decay becomes comparable to a power-law; b) DFA curves for AR(1) data; as $c \to 1$, the DFA exponent $\alpha \to 1.5$; c) DFA exponents for AR(1) data, $0 \leq c < 1$ with step 0.001. For each c, twenty samples of length 20K generated to obtain error bars; d) DFA1-propagator ζ and ACF-propagator c for AR(1) data of length 20,000, twenty samples for each c provide error bars; grey (ACF data) and black (DFA data) curves are almost identical, indicating the accuracy of the calibration.

At critical states close to the bifurcation, when indicator $c \to 1$, the slow exponential decay is well approximated (in terms of absolute deviations in the log-log plot of $C(s)$) by a power law in the short-term regime (10-100 time units), by auto-correlation function exponent $\gamma \to 0$ (see Figure 2.3a, where γ is the slope of the fitted lines). (In the long-term regime, an AR(1) process is not power-law correlated, nor is the exponential decay equivalent to the power-law decay.) Hence in the short-term regime, detrended fluctuation analysis can be used for estimation of the indicator (the same time scale for DFA). The DFA exponent α is estimated as a coefficient of linear regression in a log-log plot of $F(s)$ vs s. Figure 2.3b shows that the AR(1)-process becomes non-stationary at critical values of c, and when $c \to 1$, the DFA exponent $\alpha \to 1.5$, which corresponds to a random walk state of the system. Note, however, that this kind of non-stationarity is a necessary but not sufficient condition for bifurcation. In the case of a 'blind' test, the trend of the indicator should be considered as a warning of potential system bifurcation.

[Held and Kleinen, 2004] use this observation to argue that in a multi-dimensional system one should project onto the first Empirical Orthogonal Function (EOF), which is the basis vector for which the variance of the system is maximised. This is the one-dimensional basis in which the rise in autocorrelation (and also variance) will occur.

The rise in ACF, the "fingerprint" of the bifurcation, is detected by calculating the ACF of a segment of the time series, say $\{X_i\}_{i=m-k}^{m}$, then sliding this window by one point ($m-k+1$ to $m+1$) and calculating the ACF of this new time series segment.

Spectral properties of the time series may also be used as indicators [Kleinen et al., 2003], in particular the scaling exponent of the power spectrum. The use of spectral properties as an EWS will be useful when the tipping point is associated with a change in the structure of the noise, or the stationary system becoming non-stationary. However, a shift from a dominance of short-scale memory to long-scale memory, seen in the measurement of the power spectrum scaling exponent, will be associated with a rise in autocorrelation, and so in this respect the PS and the ACF are related. How close they are numerically is illustrated in Figure 2.5.

In order to quantify uncertainties, one can consider a range of window lengths. Consider a fixed window length of 100 points in order to compare the PS-indicator to the ACF(1)-indicator used by [Held and Kleinen, 2004] and the DFA-indicator used by [Livina and Lenton, 2007], when applied to artificial data illustrated in Figure 2.6. This experiment is in agreement with an equivalent result in [Livina et al., 2012].

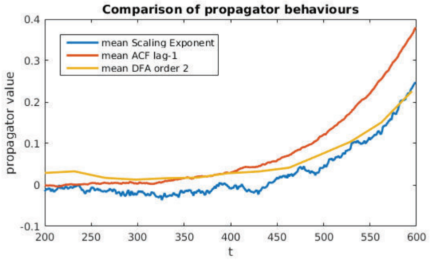

Figure 2.5: A comparison of the three autocorrelation indicators discussed. In each case the indicators are calculated for all 100 series, using a window of 200 points, and then the mean of all is taken. The scaling exponent and DFA coefficient have been scaled linearly to fit on the same axis.

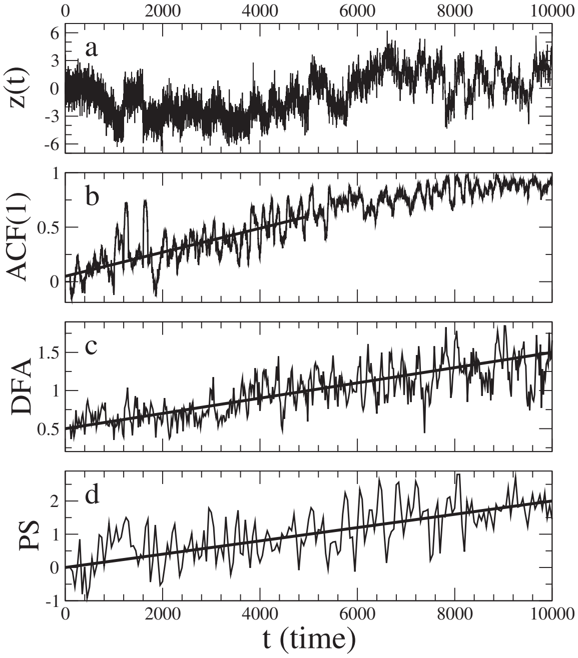

Figure 2.6: Artificial data with ACF(1)-, DFA- and PS-indicators. (a) A time series is constructed by concatenating 50 sub-series of length 200 where the scaling exponent β within each sub-series is constant and increases over the whole series from 0 (white noise) to 2. (b,c,d) The ACF(1)-, DFA- and PS-indicator methods are applied with window size 100. As the PS-indicator increases from 0 to 2, the ACF(1)-indicator increases from 0 to 1, and the DFA scaling exponent α increases from 0.5 to 1.5. Lines are added to show the linear trends; note that the ACF(1)-indicator does not increase linearly as it approaches 1.

2.2 Detecting Tipping: Potential Analysis

Detecting tipping points is performed using dynamical potential analysis. It detects a transition or a bifurcation in a series at the time when it happens, which is illustrated in innovative contour plot mapping the potential dynamics of the system [Livina et al., 2010, 2011].

We treat the climate system as a nonlinear dynamical system which can possess multiple states, with shifts between these different climate states induced by stochastic forcing. A one-dimensional conceptual model for this is given by the Langevin equation

$$dz = -U'(z)\,dt + \sigma dW, \tag{2.8}$$

where $U(z)$ is a potential function, σ is the noise level and W denotes a standard Wiener process. The state variable z represents a large-scale climate variable such as Greenland palaeotemperature, or ice-core proxy (δ^{18}O stable water isotope) records, or some other lead variable of the dynamical system.

2.2.1 Detecting the number of system states

The shape of the potential is given by a polynomial

$$U(z) = \sum_{i=1}^{L} a_i z^i, \tag{2.9}$$

where the order L is even and the leading coefficient a_L is positive for Eq. (2.8) to possess a stationary solution. The order of the polynomial controls the complexity of the potential. Increasing values of L allow more states to be accommodated; for example, a fourth-order polynomial can capture a system with two states (double-well potential) [Kwasniok, 2018, Kwasniok and Lohmann, 2009].

The number of system states is estimated by means of a polynomial fit of the probability density function of the data. Suppose the system is governed by Eq. (2.8). The corresponding Fokker-Planck equation for the probability density function $p(z,t)$

$$\partial_t p(z,t) = \partial_z[U'(z)p(z,t)] + \frac{1}{2}\sigma^2\partial_z^2 p(z,t) \tag{2.10}$$

has a stationary solution given by [Gardiner, 2004]

$$p(z) \sim \exp[-2U(z)/\sigma^2]. \tag{2.11}$$

Given this one-to-one correspondence between the potential and the stationary probability density of the system, the potential can be reconstructed from time series data of the system as

$$U = -\frac{\sigma^2}{2} \log p_{\mathrm{d}}, \tag{2.12}$$

where p_{d} is the empirical probability density of the data. This is estimated using a standard Gaussian kernel estimator [Silverman, 1986]. We applied Matlab code with built-in `kdensity` function and the NAG Fortran library (computing the Gaussian kernel density estimator using a fast Fourier transform). The estimator is

$$\hat{f}(z) = \frac{1}{nh} \sum_{i=1}^{n} K\left(\frac{z - z_i}{h}\right), \tag{2.13}$$

where K denotes the Gaussian kernel, z_i are the data points, n is the length of the data set and h is the bandwidth controlling the smoothness of the estimator. Following [Silverman, 1986], we chose $h = 1.06s/n^{1/5}$, where s is the standard deviation of the data set. Then least-square fits of $-\log p_{\mathrm{d}}$ weighted with the probability density of the data [Kwasniok and Lohmann, 2009] with polynomials of increasing even order L are calculated, starting with $L{=}2$, until a negative leading coefficient a_L is encountered. The polynomial of highest degree before first obtaining a negative leading coefficient is considered the most appropriate representation of the probability density of the time series, both locally and globally, avoiding overfitting of sampling fluctuations in the probability density.

After the occurrence of a negative leading coefficient, it may happen that a higher-order polynomial again has a positive leading coefficient. However, it appears reasonable not to consider those higher-order polynomials after the first occurrence of a negative leading coefficient.

After the initial detection, we further estimate the number of states S as

$$S = 1 + \frac{I}{2}, \tag{2.14}$$

where I is the number of inflection points of the fitted polynomial potential of appropriate degree L as described above. This definition takes into account not only the degree of the polynomial but its actual shape. We only look at even-order potentials with positive leading coefficient. These have positive curvature both at minus and plus infinity. Thus, inflection points can only occur in pairs (if any). Any potential has at least one state (with no inflection points). Then we count one further state for each pair of inflection points. This can be either a real minimum (well) or just a flattening in the potential (corresponding to degeneracies in the potential [Kwasniok and Lohmann, 2012]; definition (2.14) systematically accommodates both possibilities. The number of inflection points is numerically given as the number of sign changes in the second derivative on a fine enough mesh.

A simplified version of definition (2.14) would be to only count the relative minima (real wells) in the potential. We opt here for the more comprehensive definition of system states as degenerate potentials have been shown to occur in the context of ice-core records [Kwasniok, 2018, Kwasniok and Lohmann, 2009].

Examples of the polynomial fits for simulated data are given in Figure 2.7. Given double-well potential data ($U(z){=}z^4{-}2z^2$), the probability density of the time series has a bimodal shape. The empirical potential given as $U{=}{-}\frac{\sigma^2}{2}\log p_{\mathrm{d}}$ can be fitted by polynomials of various orders, but only those of degree 2 and 4 have a positive leading coefficient. Higher-order polynomials have negative leading coefficients — although fitting the data better and better locally, they are clearly inadequate globally. The fourth-order polynomial is the appropriate representation of the probability density and the method identifies the time series as double-well potential.

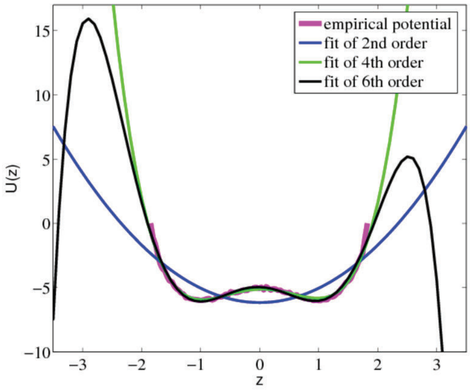

Figure 2.7: Empirical potential (magenta curve) of double-well potential artifical data and its polynomial fits of 2nd, 4th and 6th orders. The fit with 2nd order polynomial is not as accurate as with the 4th order, whereas the fit of 6th order approximates the potential only locally, becoming negative for large $|z|$.

To plot the results with the estimated number of system states, we use a special contour plot, where the x-axis is the time scale of the record, the y-axis is the length of the sliding window moving along the record, and the color denotes the number of detected states, mapped into the middle of the sliding window in the x-axis. Here, red colour denotes one detected state, green — two, cyan — three, purple — four. Note that there is an ambiguity in mapping the results (beginning, middle or end of the windows), and since we consider data aggregated for kernel distribution in each window, it is logical to map the result into the middle.

This 2D contour plot can be used to help detect bifurcations in the time series. A change in the number of states (vertical change of colour) along all or most time scales (window sizes) indicates the appearance or disappearance of a state (or states). This is not precisely equivalent to a bifurcation, which occurs when a marginally stable state becomes degenerate or a degenerate state becomes marginally stable (i.e. the potential exhibits a plateau or point of inflection). Instead, when the colour change indicates a decline in the number of states, then one can infer that a bifurcation occurred sometime beforehand. Conversely, when the colour change indicates an increase in the number of states, then one can infer that a bifurcation may be approaching. Where e.g. forcing is strong, the change in the number of states may be closely associated in time with a bifurcation, but there is no guarantee of this. To characterize the changing stability of states and thus help locate bifurcations, we must estimate the potential shape.

The sliding windows should be larger than the time scale of jumps between the different states for the system to be likely to have explored the whole phase space (visited all existing states) in that time frame. When the sliding windows are too small, they may detect a smaller number of states than in the global potential, due to the temporal scale restriction.

The described procedure is used to estimate the number of the system states only. To calculate the potential coefficients, we use the unscented Kalman filter (UKF, see section 2.2.2), which is an independent method. For instance, consider a degenerate potential, where two wells bifurcate into one; the method estimating the number of states will correctly identify it as one-well-potential, but since the potential curve may still be asymmetric, it may still require four coefficients (some of them equal to zero) to be described accurately.

We opt here for the more comprehensive definition of system states as degenerate or nearly degenerate potentials have been shown to occur in the context of ice-core records [Kwasniok, 2018, Kwasniok and Lohmann, 2009]. Moreover, even if the system possesses real potential wells, changes in the number of states (bifurcations) tend to be picked up more quickly when using the inflection point definition as this is a more sensitive feature.

To test the method for estimating the number of system states, we generated four sets of artificial data for four different potentials: one-well ($U(z) = z^2$), double-well ($U(z) = z^4 - 2z^2$), triple-well ($U(z) = z^6 - 4.5z^4 + 5z^2$) and four-well potential ($U(z) = z^8 - 6.5z^6 + 13z^4 - 8z^2$) (Figure 2.8a). Simulated data of these systems were obtained by numerically integrating eq. (2.8) using the Euler scheme. The noise level $\sigma = 1.5$ was adopted, resulting in a mean state transition time (Kramers waiting time) similar to that occurring in the palaeoclimatic records when interpreting system time units as kyr. The sampling interval is 0.05 system time units. For each of the potentials, 150 time units worth of data were generated, resulting in 3000 data points. The subsets are combined to produce a record with 12 000 data points, corresponding to time running from 0 to 600 (Figure 2.8b).

In the artificial data, the number of states is generally identified correctly; changes are picked up quickly and reliably (Figure 2.8c). For the smallest time windows, there are sporadic misidentifications of the number of states, due to poor statistics and occasional changes of dynamics in small subsets. Generally time windows of order 400 data points (20 time units) are sufficient to get reliable results.

The potential contour plot visualises bifurcations according to the number of detected system states. It identifies bifurcations as the change in the number of states along all time scales (y-axis shows the size of the sliding window with subset of data where the number of states is assessed). If no such pattern is observed, there is no bifurcation in the series. If the potential plot is “patchy”, indicating constantly varying potential of the system, it can be classified as non-potential, i.e. no global deterministic quasi-stationary potential exists to describe the dynamics of the system.

[Ditlevsen and Johnsen, 2010] argue that between 25 and 45 kyr BP, the Greenland paleotemperature did not experience a system bifurcation but rather oscillated between stable system states. It is an interesting confirmation of the result of [Livina et al., 2010], where by using a different technique, the potential contour plot, it was shown that in the mentioned time interval the system potential has stable double-well structure. The same can be concluded using DFA indicator, but not the ACF-indicator, which does rise to critical value 1, together with variance indicator, and thus provides false alarm of a tipping point that does not happen, see Figure 2.9.

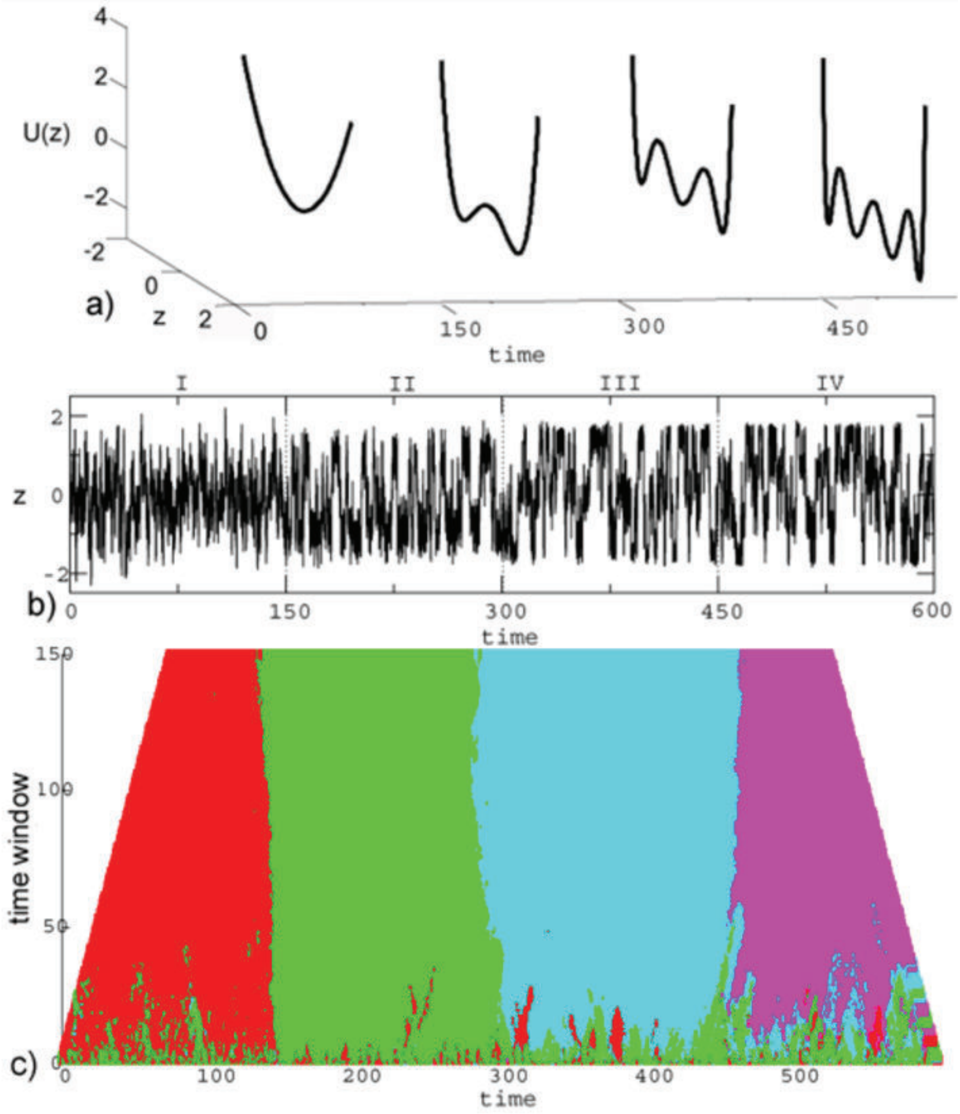

Figure 2.8: Detection of the number of states in artificial data: **(a)** four different potentials with increasing number of wells from one to four; **(b)** the resulting time series; **(c)** contour plot of the number of detected states versus time and size of sliding window (red – one well, green – two, cyan – three, purple – four). The number of detected wells is mapped at the end of the sliding time windows.

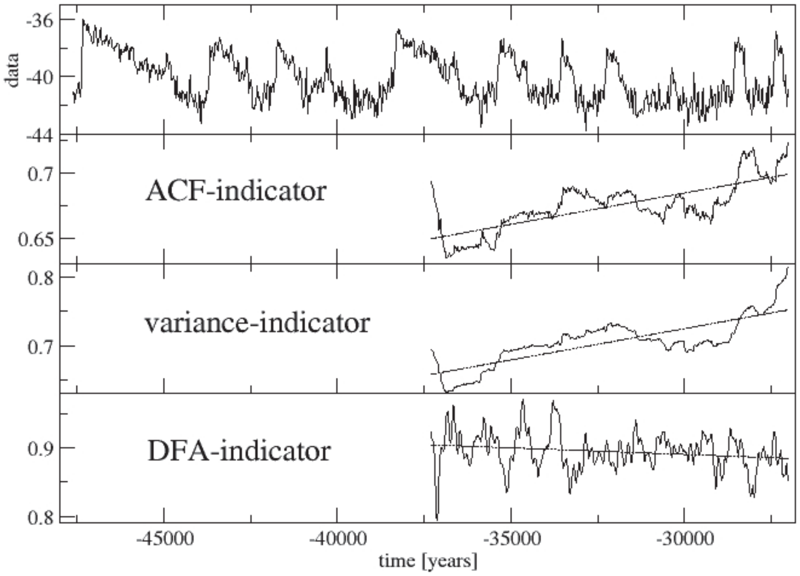

Figure 2.9: GRIP paleotemperature proxy in time interval 45-25 kyr BP, where there is no system bifurcation, which is confirmed by DFA-indicator, whereas ACF-indicator and variance give false alarm of tipping. Dashed lines show the slopes of linear regressions, with indication of early warning signal in ACF and variance, but not in DFA-indicator. The result is confirmed by different techniques in [Cimatoribus et al., 2013, Livina et al., 2010].

As introduced in [Livina et al., 2010], potential analysis allows one to detect system bifurcations with accurate timing of structural changes of the system potential. It is a powerful visualisation tool for tipping and failure detection. For specific types of failures in dynamical systems, new states may be appearing or old states vanishing, and potential analysis is helpful in their identification.

The polynomial fitting of the probability density in sliding windows of varying size provides a quick dynamical portrait of the time series, which is not achievable by mere calculation of potential coefficients over multiple subsets of time series. The potential contour plot is a visualisation of the dynamics (often hidden in time series and not visible by eye) at different time scales. It is possible to run the UKF algorithm for a fixed high-order polynomial, calculating all coefficients and then showing which of them are practically zero and thus estimating the actual structure of the potential. However, to analyse this information, one needs a visualization, and this is where the potential contour plot serves the purpose. Moreover, when the statistics are poor (less than 700–800 datapoints), the polynomial fitting is still capable of detecting the number of states correctly (see Figure 2.10).

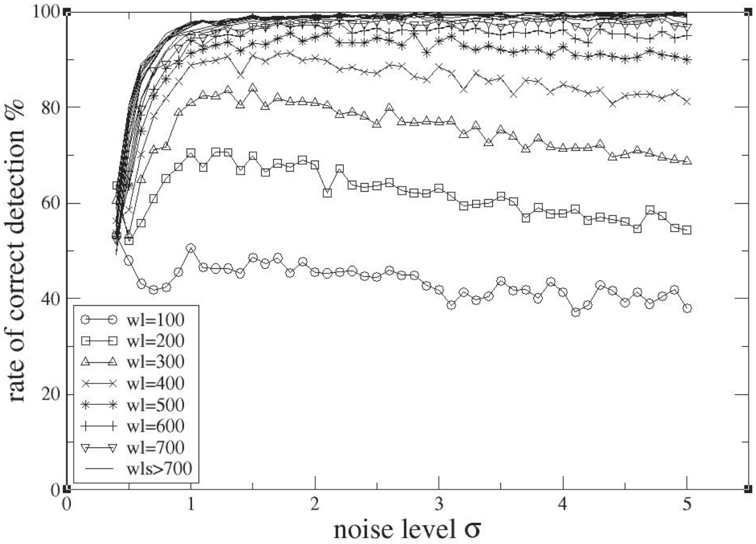

Figure 2.10: The rate of detection of two wells in artificial double-well-potential data, $U(z) = z^4 - 2z^2$, with noise level σ values between 0.4 and 5 (step 0.1, 1000 samples of length 10K for each value of σ). Detection is performed in windows of increasing size, from 100 to 2000 with step 100. The two potential wells are of equal depths 1, so for small values of σ the potential barrier may be not crossed, and the data in small windows may remain one-well-potential locally — this explain poorer performance of the method at small noise level. For windows bigger than 700 datapoints the correct detection rate is higher than 95% for noise level above 1.5, as one can see that all the curves coincide close to 100% detection. With increasing noise level σ, for small windows the detection rate slightly declines, but still remains higher than 95% for windows bigger than 700 datapoints.

This simple approximate approach works with remarkable accuracy for data subsets of length as short as 400 to 500 data points, demonstrating above 90% rate of accurate detection, as was shown in an experiment with artificial data. For data subsets of length above 1000 points it correctly detects the structure of the potential with a rate of 98% [Livina et al., 2011].

2.2.2 Estimating potential coefficients using the unscented Kalman filter

The estimation of the number of system states is augmented by a nonlinear dynamical estimation of the shape of the potential. The numerical method to derive the coefficients of the potential uses the unscented Kalman filter [Julier et al., 2000, Kwasniok, 2018, Kwasniok and Lohmann, 2009, Sitz et al., 2002]. The method incorporates measurement noise and model error or inadequacy.

The UKF is a nonlinear extension of the conventional Kalman filter. It allows for recursive estimation of unobserved states and parameters in both deterministic and stochastic nonlinear models from incomplete, indirect, and noisy observations. Unlike the extended Kalman filter, the UKF does not linearize the system dynamics but keeps its full nonlinearity. It consistently propagates the first and second moments of the state and parameter estimates.

The dynamical evolution of the climate variable is given by a discretization of Eq. (2.8) using the Euler scheme with step size h:

$$z_t = z_{t-1} - hU'(z_{t-1}) + \sqrt{h}\sigma\eta_t \tag{2.15}$$

The observation equation is

$$y_t = z_t + \epsilon_t. \tag{2.16}$$

y_t is the observed variable which is here identified with the climate records; the state variable z_t is not directly observable. The Gaussian observational noise ϵ_t with zero mean and variance R captures measurement uncertainty and/or model error. The UKF simultaneously estimates the unobserved state z_t and the potential parameters $a_1, \ldots, a_L$ in Eq. 2.9 from only a time series of the noisy observations y_t. An augmented state vector $\mathbf{x}$ of dimension $n = L+1$ is created by merging the state variable and the parameters: $\mathbf{x} = (z, a_1, \ldots, a_L)$. The dynamical evolution of $\mathbf{x}$ is given by Eq. (2.15) augmented by a constant dynamics for the parameters.

Let $\hat{\mathbf{x}}_{t-1|t-1}$ be the estimate of the augmented state vector at time $t-1$ having processed all data up to time $t-1$ and $\mathbf{P}^{xx}_{t-1|t-1}$ its covariance matrix. The probability density of the augmented state vector is represented by $2n$ carefully chosen so-called sigma points [Julier et al., 2000]. The sigma points are propagated through the augmented dynamical equation, leading to transformed means and covariances $\hat{\mathbf{x}}_{t|t-1}$, $\hat{y}_{t|t-1}$, $\mathbf{P}^{xx}_{t|t-1}$, $\mathbf{P}^{xy}_{t|t-1}$ and $P^{yy}_{t|t-1}$. When reaching a new observation y_t, the estimates of the state and the parameters are updated using the ordinary Kalman update equations

$$\begin{aligned} \hat{\mathbf{x}}_{t|t} &= \hat{\mathbf{x}}_{t|t-1} + \mathbf{K}_t(y_t - \hat{y}_{t|t-1}) && (2.17) \\ \mathbf{P}^{xx}_{t|t} &= \mathbf{P}^{xx}_{t|t-1} - \mathbf{K}_t P^{yy}_{t|t-1} \mathbf{K}^{\mathrm{T}}_t && (2.18) \end{aligned}$$

where $\mathbf{K}_t$ is the Kalman gain matrix given by

$$\mathbf{K}_t = \mathbf{P}^{xy}_{t|t-1} (P^{yy}_{t|t-1})^{-1}. \tag{2.19}$$

See [Kwasniok and Lohmann, 2009] for more technical details on the estimation of potentials with the UKF.

2.2.3 Estimating the noise level

Finally, accurate estimation of the noise level is crucial for the determination of the potential coefficients. To this end, we first apply wavelet denoising of the series with soft thresholding and Daubechies wavelets of 4th order. By subtracting this series from the initial record, we obtain a time series that is used for estimation of the dynamical noise level σ. Using this series, the UKF algorithm is applied in a set of runs with a range of noise levels centered at the roughly estimated value of σ. Then for the coefficients obtained by the UKF algorithm, the probability distribution is calculated, and the most appropriate value of the noise level is chosen by comparison of cumulative distributions of models and data Φ_m and Φ_d (the latter is the integral of the empirical probability distribution p_d) at the minimum value of measure

$$D = \max_z |\Phi_m(z) - \Phi_d(z)| \,. \tag{2.20}$$

This measure allows one to find the value of the noise level that provides the minimal global discrepancy between modelled and observed cumulative distributions, i.e. to find empirically the most plausible noise level given the time series of the system.

The noise is separated only for the purpose of estimating the noise level σ, which is further used for calculation of the potential coefficients. The noise is never separated when we calculate the potential contour plot. Thus, we apply the separation of noise to eliminate the influence of nonstationarities on the estimation of noise level.

2.3 Forecasting Tipping

Forecasting time series, especially those experiencing transitions or bifurcations, is a major challenge in time series analysis. Various techniques have been developed, which derive statistical properties of a time series prior to the moment of forecast to simulate these properties in the ”future” — this may be done to the historic data and is called ”hindcast” (forecast in the past). Hindcast is a good technique for testing forecast accuracy.

However, because tipping events may vary widely (bifurcational or transitional, abrupt or gradual), in many cases forecasting these changes are often inaccurate; in many cases instead of a single forecast one performs an ensemble forecast with multiple trajectories, which also provide uncertainty quantification of the ensemble.

In general, time series forecast techniques rarely achieve a long horizon of accurate prediction, by ”long” considering the time scale that includes a nonstationary transformations of interest. Depending on sampling rate, this horizon may vary between hundreds and thousand points, and by taking into account several statistical properties of the historic data, such as combination of trends and stochastics components, sometimes it is possible to achieve accurate forecasts with horizon of several hundred datapoints even for nonstationary data [Livina et al., 2013].

Potential forecasting was introduced in [Livina et al., 2013] and is based on extrapolation of the dynamics of probability density with simultaneous reconstruction of temporal scaling (long-term autocorrelations) from the historic data. The horizon of such forecasts may include the nonstationary part of a time series, and its performance may be assessed using various error measures. This technique is applied to artificial data to compare its performance with other time series forecasting techniques. However, for successful potential forecasting, the dynamical system should exhibit potential behaviour, i.e. there should exist a stochastic potential equation that would adequately model the time series. This is not always the case in real systems, as was shown in [Livina et al., 2012].

The technique of potential forecasting uses polynomial coefficients of the orthogonal approximation of the empirical probability distribution. After reconstruction of the system kernel distribution, a time series is generated using a rejection sampling technique, and then the obtained dataset is sorted according to the initial data in order to reconstruct the temporal correlations in the series.

For deterministic models with a linear perturbation of the system potential, it is possible to analyse the proximity to a bifurcation [Surovyatkina et al., 2005], but the problem of the real-world systems is that the potential may be completely unknown, and the perturbation may not be linear. In fact, all terms of the polynomial potential may vary simultaneously. Such complex dynamics means that a possible way of analysing the system is to derive the sequence of the potential coefficients and attempt to identify the leading parameter that perturbs the system potential.

The potential can be reconstructed from time series data of the system as

$$U(x) = -\frac{\sigma^2}{2}\log p_{\mathrm{d}}(x), \tag{2.21}$$

which means that the empirical probability density p_{d} has the number of modes corresponding to the number of wells of the potential.

2.3.1 Approximation of the probability density

To approximate the empirical probability density, we use the orthogonal (in the interval $[-1, 1]$) Chebyshev polynomials of the first kind:

$$\begin{aligned} &T_0(x) = 1, \\ &T_1(x) = x, \\ &\ldots \\ &T_{n+1}(x) = 2xT_n(x) - T_{n-1}(x). \end{aligned} \tag{2.22}$$

The polynomial $T_n(x)$ has n zeros in the interval $[-1, 1]$ at points

$$x = \cos\left(\frac{\pi(k-1/2)}{n}\right), \quad k = 1, 2, \ \ldots, \ n.$$

The approximation of a function $f(x)$ can be done by using a truncated (finite N) sum of the following form

$$f(x) \cong \left(\sum_{k=0}^{N-1} c_k T_k(x)\right) - \frac{1}{2}c_0, \tag{2.23}$$

where c_k are the coefficients obtained by discrete cosine transform of the vector of nonuniformly spaced samples of the considered function over the sampling grid

$$x_k = \cos\left(\frac{\pi(k-1/2)}{n}\right), \quad k = 1, 2, \ \ldots, n.$$

For an arbitrary interval $[a, b]$ it is necessary to transform variables as

$$y = \frac{x - 0.5 \cdot (b + a)}{0.5 \cdot (b - a)}.$$

A good example of the above calculations is given in [Neagoe, 1990].

When decomposition (2.23) is obtained according to the particular time series problem to be analysed, the resulting polynomial is expanded, thus producing the final coefficients

$$f(x) \cong \left(\sum_{k=0}^{N-1} \tilde{C}_k x^k \right). \tag{2.24}$$

2.3.2 Linear extrapolation of the coefficients and forecast time series

When the approximating polynomial is derived, the decomposition coefficients $\tilde{C}_k$ (Eq. 2.24) are linearly extrapolated using a set of preceding values. The intervals of these, as well as the extension of the extrapolated interval can be chosen according to the particular time series to be analysed.

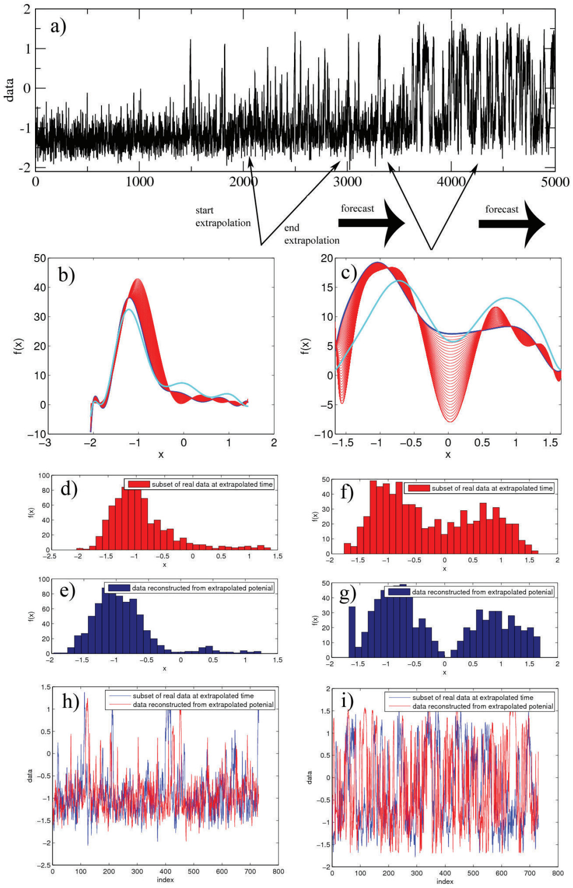

Figure 2.11: Artificial data bifurcating from one-well to double-well dynamics — two hindcasts are demonstrated. Panels: a) time series – the dynamics is potentially extrapolated in two intervals: from 2100 to 2900 and from 3400 to 4200; b) emprical probability density (Chebyshev-polynomial approximations) for the interval from 2100 to 2900: blue curve is the initial probability density, cyan is the final curve; red curves are extrapolated at equal steps; c) the same as (b) for the interval from 3400 to 4200; (d) and (e) are histograms for extrapolation at point 2900: comparison of the actual histogram and the result of extrapolation; (f) and (g) are the same as (d) and (e) for the point 4200; h) comparison of the forecast and actual data in interval from 2900 to 3630: i) comparison of the forecast and actual data in interval from 4200 to 4930.

The initial pdfs (those providing sequences of potential coefficients to be extrapolated) are estimated for the fixed time interval, in a sequence of subsets prior to the forecast starting point (in sliding windows). The extrapolation of parameters provides adequate scaling of the pdf, and no further normalisation is necessary. Figure 2.11 demonstrates application of the potential forecast on artificial data of a dynamical system with a gradual bifurcation from one to two system wells.

Once a new probability density is calculated, we generate a forecast time series using a rejection sampling algorithm (see, for instance, [Gilks and Wild, 1992]). This provides an artificial series with the prescribed distribution, but this may be not enough for obtaining a realistic forecast time series, because the ordering of the series (and hence scaling properties like long-term memory) should be reconstructed according to the initial data. For this purpose, we apply so-called ”sorting” of time series, that means arranging its values in the same order as in the initial data (before the forecast started), thus reproducing realistic correlations (because their distributions are already very similar due to the extrapolation of probability density).

Sorting is a simple numerical algorithm which uses ranking of the values of two series, initial subsample and forecast subsample. However, this should be done with care, especially in data with seasonality: if there is a seasonal trend, it is very important to sort the forecast series according to observed data of the same interval of the year, to reproduce seasonal patterns on the day. This is achieved by going back along the series with step equal to the seasonality period (365 for daily data or 12 for monthly data). Since certain years may be anomalous in fluctuations (due to internal variability in the system), the initial data used for sorting may be an average over several years starting from the same date in a year (for instance, March 1st in several consecutive years). This average is then used to sort the forecast series starting on March 1st and projecting into the future. Application of the algorithm on real data, Arctic sea ice area, is shown in Figure 2.12.

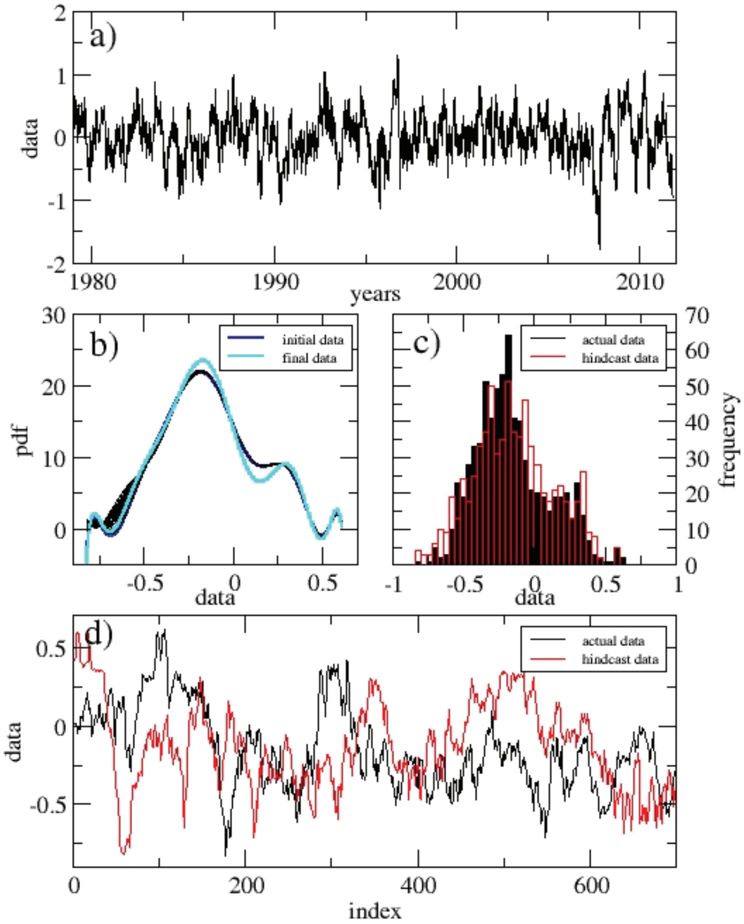

Figure 2.12: Arctic sea-ice area fluctuations after deseasonalising and removal of quadratic decreasing trend: a) time series; b) hindcast empirical probability density (Chebyshev-polynomial approximations), where blue curve is the initial statistics at the beginning of the hindcast, black curves are extrapolated densities up to 100 time units ahead, cyan curve is the real pdf at the end of the forecast, for comparison with extrapolation; c) histograms of the forecast and real data at the end of extrapolation; d) time series corresponding to histograms in the panel c): 2-year-long hindcast.

2.3.3 Uncertainties and applicability; criteria of performance

It is necessary to note that minor uncertainties are introduced at various steps of the forecasting algorithm: first when potential stochastic model is used as an approximation of dynamics, then when the empirical probability density is approximated by Chebyshev polynomials (minor outliers); furthermore, the polynomial coefficients are linearly extrapolated, which means that the actual dynamics is most correctly forecast only in the case of linear evolution of such coefficients.

In many cases of abrupt highly nonlinear dynamics the linear extrapolation of the decomposition coefficients may produce the empirical distribution with large deviations, especially in the case of non-stationarity of the data. Bootstrapping of the decomposition of coefficients can be applied according to [Mudelsee, 2012]. Based on bootstrapping techniques, it is possible to consider blocks of data in a chosen subset x of size

$$L = NINT\left(W^{1/3}\frac{\sqrt{6}a_1(x-\bar{x})}{1-(a_1(x-\bar{x}))^2}\right), \tag{2.25}$$

where $NINT(\cdot)$ is the nearest integer function, W is the window length, a_1 is the lag-1 autocorrelation, $\bar{x}$ is the mean value of the subset x. This block length selector was derived in [Mudelsee, 2012] from [Sherman et al., 1998], who adapted a formula from [Carlstein, 1986]. For $a_1 \rightarrow 0$, L is chosen equal to 1; when the denominator of Eq. (2.25) tends to 0, L is chosen equal to $n-1$.

In the case of nonstationary data, when the probability distribution varies within the data subset, bootstrapping provides estimates of the partial probability distributions, which may deviate from the average quite significantly. In practical terms, applying bootstrapping for estimation of the decomposition coefficients in non-stationary data provided worse results in the considered samples, with the forecast time series of poorer accuracy than the single-estimated probability density functions. A possible solution to this could be modification of the bootstraping algorithm, where instead of mean value removal a more sophisticated detrending is applied.

Furthermore, the important parameter that affects the skill of the forecast is the extrapolation period. The skill of the forecast drops with increase of its size.

Obviously, in the case of abrupt changes using linear extrapolation of coefficients may prove unsatisfactory. The best results are obtained when the data undergoes gradual dynamic change and the forecast horizon is within 100 time units (which means 3 months for daily data and up to 8 years for monthly data).

Forecast validation

The following metrics can be used to analyse the goodness of fit of forecasting models to compare observed and modelled time series. If x_o and x_m of length n are the observed and modelled series respectively, $i = 1, \ldots, n$ is the time index and $\bar{x_o}$ is the mean value of the observed time series, then:

Nash-Sutcliffe (NS) efficiency:

$$1 - \left(\frac{\sum_{i=1}^{n} ({x^i}_m - {x^i}_o)^2}{\sum_{i=1}^{n} ({x^i}_o - \bar{x_o})^2} \right); \tag{2.26}$$

Root Mean Squared Error (RMSE):

$$\sqrt{\frac{\sum_{i=1}^{n} (x_m - x_o)}{n}}; \tag{2.27}$$

Percent Bias (PB):

$$\left[\frac{\sum_{i=1}^{n} ({x^i}_m - {x^i}_o)}{\sum_{i=1}^{n} {x^i}_o} \right] \times 100. \tag{2.28}$$

If the observed and modelled series are identical, then $NS = 1$, $RMSE = 0$, and $PB = 0$, a deviation from these values indicate a pointwise difference between the observed and modelled time series.

Chapter 3

Multivariate tipping points

Early warning signal techniques are frequently applied to univariate data, often from a single data source or model output. Multivariate data may be from multiple observations of the same variable or several different variables measured in the same location, or a combination of different variables at multiple locations. In the first case it may be sufficient to attempt an EWS detection on each time series individually, but it may be not possible to detect an EWS in $X(t)$ nor $Y(t)$ but, when considered together, they may produce an EWS.

It may therefore be worthwhile to investigate the possibility of a two (or more) dimensional analogue to existing one-dimensional techniques. In the case of gridded data it is common to use Empirical Orthogonal Functions, also known as Principal Component Analysis (PCA), to reduce the dimensionality [Jolliffe, 1986, von Storch and Zwiers, 2002], this technique has also been more specifically applied to EWS methods [Held and Kleinen, 2004, Kwasniok, 2018].

3.1 Multivariate extension of the autocorrelation function

In [Williamson and Lenton, 2015], the authors introduced a method to anticipate bifurcations in time series generated from a known analytical system. The method is intended as a higher-dimensional analogue of the one-dimensional ACF(1) indicator. There is equivalence between the equation used to calculate the lag-1 autocorrelation coefficient a in the one-dimensional time series $\{x_t\}$:

$$a = \frac{\sum_t (x_{t+1}x_t) - \overline{x}^2}{\sum_t (x_t^2) - \overline{x}^2}, \tag{3.1}$$

and the equation given by the authors:

$$A = \left[\sum_t (\mathbf{x}_{t+1}\mathbf{x}_t^\top) - \overline{\mathbf{x}}^2\right] \left[\sum_t (\mathbf{x}_t\mathbf{x}_t^\top) - \overline{\mathbf{x}}^2\right]^{-1} \tag{3.2}$$

where $\{\mathbf{x}_t\}$ is a multivariate time series. One may model the system equations of the time series $\{x_t\}$ as the auto regressive system

$$x_t = ax_{t-1} + c + \varepsilon_t, \tag{3.3}$$

where c is a scalar value and ε_t is Gaussian noise, using the lag-1 autocorrelation coefficient a given by equation (3.1). In the same way, the matrix A may be used to model a multivariate system using the analogous form

$$\mathbf{x}_t = A\mathbf{x}_{t-1} + \mathbf{c} + \varepsilon_t, \tag{3.4}$$

where $\mathbf{c}$ is a constant vector and ε_t is a vector with each element being independent Gaussian white noise. In this system a tipping point occurs when an eigenvalue of A has the value one. It is therefore changes in the eigenvalues of A that must be studied. It is proposed in [Williamson and Lenton, 2015] that the eigenvalues of the Jacobian of the system equations, evaluated at a stable point $\mathbf{x}_\star$, $J(\mathbf{x}_\star)$, should be studied because by studying the system equations analytically one can see the effect that changing parameters will have on the location of $\mathbf{x}_\star$ and therefore the eigenvalues of the Jacobian. However, rather than attempting to calculate the Jacobian matrix of a dynamical system from a time series, the authors calculate A directly using equation 3.2, and use the approximation

$$A = \mathbf{I} + J(\mathbf{x}_\star)\Delta t \approx \exp(J(\mathbf{x}_\star)\Delta t) \tag{3.5}$$

to recover the Jacobian. The Jacobian eigenvalues λ_k are therefore given by the relations

$$\begin{aligned} \mathbb{R}(\lambda_k) &= \frac{1}{\Delta t} \ln |a_k| \\ \mathbb{I}(\lambda_k) &= \frac{1}{\Delta t} \phi_k, \end{aligned} \tag{3.6}$$

where a_k and ϕ_k are the magnitudes and phases of the complex eigenvalues of A. Calculating these eigenvalues of the matrix A, thus locating a tipping point in the autoregressive model in equation 3.4, is approximately analogous to finding the eigenvalues λ_k of the Jacobian matrix J, which gives information about the system at the stable point. The authors of [Williamson and Lenton, 2015] use this analogy to justify the use of the multivariate autocorrelation as a tipping point indicator.

The method then involves calculating the matrix A from the given time series data in a sliding window, and recovering the Jacobian eigenvalues using the equations (3.6).

3.1.1 Homoclinic bifurcation

The system given in [Williamson and Lenton, 2015] is defined by equations

$$\begin{aligned} \dot{x} &= y + \varepsilon^{(x)} \\ \dot{y} &= \alpha - x^2 + \varepsilon^{(y)} \end{aligned} \tag{3.7}$$

where $\varepsilon^{(x)}$, $\varepsilon^{(y)}$ are white noise terms. If these noise terms are ignored, the system has steady state solutions at $(x, y) = (\pm\sqrt{\alpha}, 0)$. The Jacobian of the system is given by

$$J(x, y) = \begin{pmatrix} 0 & 1 \\ -2x & 0 \end{pmatrix} \tag{3.8}$$

and has eigenvalues $\lambda = \pm\sqrt{-2x}$. The eigenvalues at the stable solutions are $\lambda = \pm(4\alpha)^{1/4}$ at $(-\sqrt{\alpha}, 0)$, indicating that this is a saddle, and $\lambda = \pm i(4\alpha)^{1/4}$ at $(+\sqrt{\alpha}, 0)$, indicating that this is a stable centre. At $\alpha = 0$ the two stable points collide resulting in a homoclinic bifurcation and there are no stable solutions for $\alpha < 0$.

The authors considered a point at or close to the stable centre $(+\sqrt{\alpha}, 0)$, with noise terms $\varepsilon^{(x)}$, $\varepsilon^{(y)}$ in equation 3.7 having standard deviation of 0.01. At this centre the eigenvalues of the Jacobian have zero real part but imaginary part of $\pm(4\alpha)^{1/4}$. If the eigenvalue with positive imaginary part is considered, it decreases as α decreases to zero. This is the early warning signal of the bifurcation.

By using a varying parameter $\alpha(t) = 0.5 - 0.005t$ in equations (3.7), with both noise terms independent with zero mean and standard deviation 0.01, we observe a bifurcation at $\alpha = 0$ and $t = 100$. The system is integrated from $t = 0$ to $t = 100$ and the solution given at intervals of $\Delta t = 0.5$.

The autocorrelation matrix A is calculated using equation (3.4) in a sliding window of 100 points. The Jacobian eigenvalues are recovered using equations (3.6) and the real and imaginary parts of the first eigenvalue (the eigenvalue with positive imaginary part) are used to analyse multidimensional early warning signals. The decreasing imaginary part of the eigenvalue is a precursor of the bifurcation, as predicted by the analysis, but it is not clear that this could serve as an EWS without the analysis of the system beforehand. For an n-dimensional system, there are $2n$ variables to analyse (n eigenvalues, each with real and imaginary parts). However, studying the principal eigenvalue, rather than the n individual components of the system, does at least incorporate part of the information of each variable into a single statistic, even if information is lost by ignoring the other eigenvalues.

3.1.2 Hopf bifurcation

Consider the system given by the polar equations [Williamson and Lenton, 2015]:

$$\begin{aligned} \dot{r} &= \mu r - r^3 \\ \dot{\theta} &= 1 + r^2. \end{aligned} \tag{3.9}$$

The system has stable solution at $r = 0$ for $\mu < 0$, which becomes unstable when $\mu \geq 0$. The nature of the bifurcation can be studied using the eigenvalues of the Jacobian of the system equations. In Cartesian rather than polar coordinates,

$$J(x, y) = \begin{pmatrix} \mu - (x+y)^2 - 2x^2 & -1 - (x+y)^2 - 2y^2 \\ 1 + (x-y)^2 + 2x^2 & \mu - (x-y)^2 - 2y^2 \end{pmatrix}. \tag{3.10}$$

At the stable point $r = 0$ (i.e. $(x, y) = (0, 0)$), the matrix $J(x, y)$ has complex conjugate eigenvalues $\lambda = \mu \pm i$. The system therefore experiences a Hopf bifurcation as μ approaches zero from below, characterised by the Jacobian eigenvalues moving from the negative-real to the positive-real half of the complex plane.

Using a varying parameter $\mu(t) = 0.05t - 2.8$, we integrate the system from $t = 0$ to $t = 60$ with the solution given at intervals $\Delta t = 0.5$. The bifurcation at $\mu = 0$ occurs at $t = 58$. Gaussian noise terms $\varepsilon^{(r)}$ and $\varepsilon^{(\theta)}$ are also added to the system equations for r and θ respectively, equations (3.9), both with zero mean and standard deviation 0.01. The result is shown in Figure 3.1.

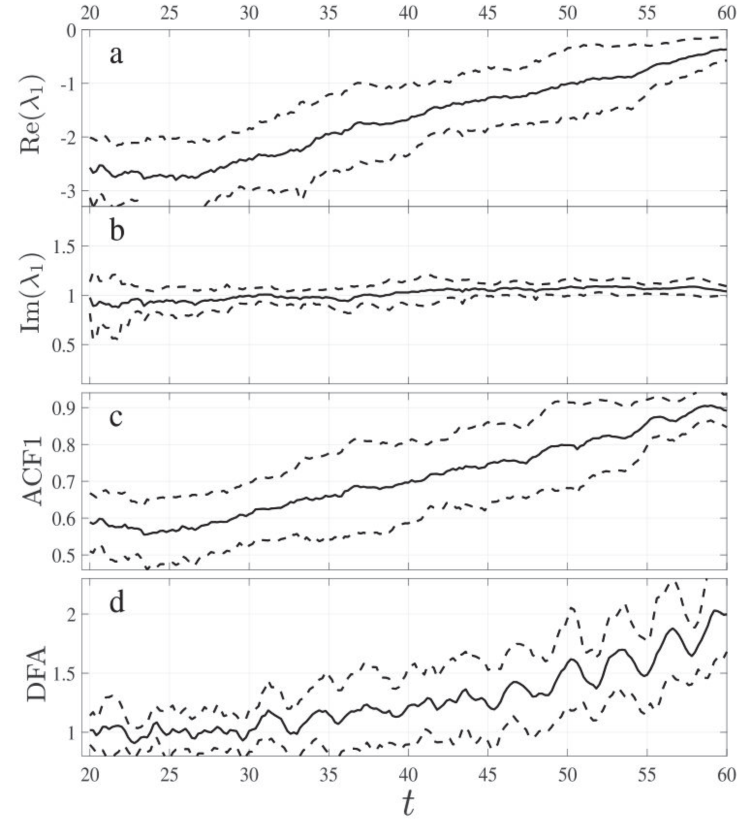

Figure 3.1: Tipping point indicators applied to a system which experiences a Hopf bifurcation at t = 58. Panels (a) and (b) show the real and imaginary parts of the first reconstructed eigenvalue of the Jacobian matrix. Panels (c) and (d) show the ACF and DFA indicators calculated with a window size of 100 points. The system was integrated ten times and the mean over ten data sets is plotted, along with error of one standard deviation (dashed lines).

As in the previous example, the autocorrelation matrix A is calculated using equation (3.4) in a sliding window of 100 points. The Jacobian eigenvalues are recovered using equations (3.6) and the real and imaginary parts of the first eigenvalue are plotted in Figure 3.1.

3.1.3 Van der Pol oscillator

[Williamson and Lenton, 2015] applied EWS to two dynamical systems, one experiencing a Hopf bifurcation and one experiencing a homoclinic bifurcation. The method can also be applied to a Van der Pol oscillator experiencing a Hopf bifurcation, as an independent test of the method.

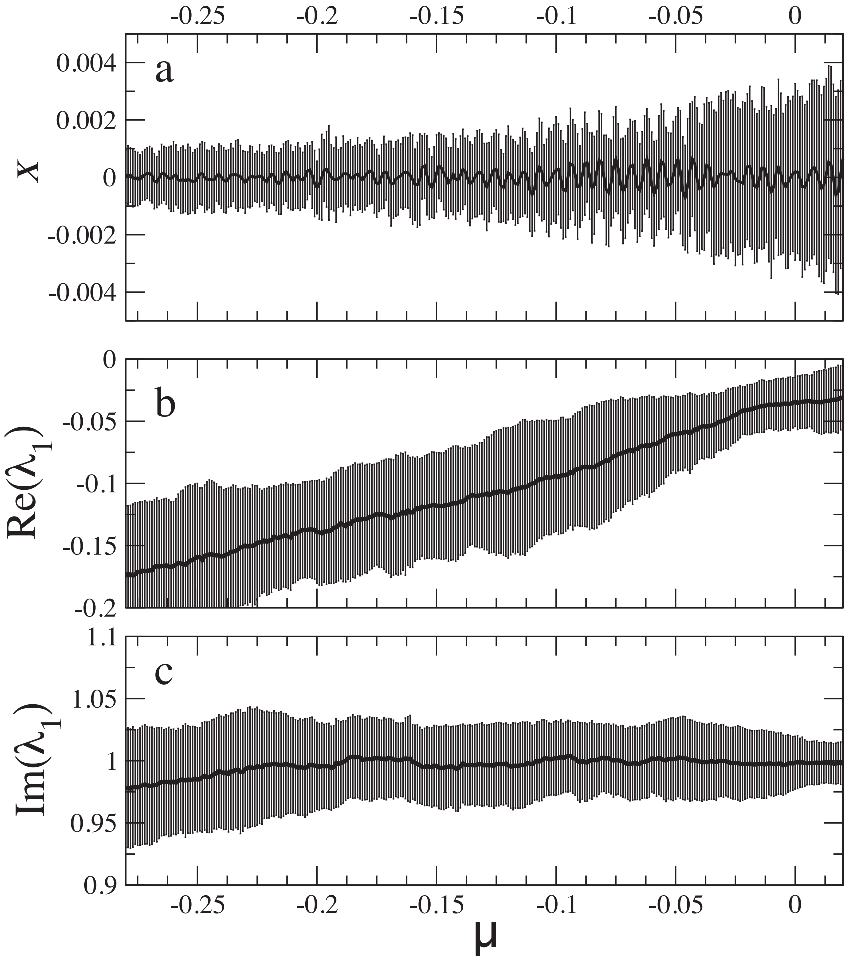

Figure 3.2: The Van der Pol system described in equation 3.11. (a) The mean of 40 separate trials is shown with error bars of one standard deviation. (b,c) The mean of the first Jacobian eigenvalues are also shown plotted against μ, with error bars of one standard deviation. An Hopf bifurcation occurs at $\mu = 0$.

Consider the Van der Pol oscillator

$$\ddot{x} - \mu(x - x^2)\dot{x} + x = \varepsilon, \tag{3.11}$$

where the stochastic forcing term ε is white noise with standard deviation 0.01. Equation (3.11) can be re-written as a coupled system of first order ODEs:

$$\begin{array}{rcl} \dot{x} & = & \mu\left(x - \frac{1}{3}x^3\right) + y \\ \dot{y} & = & -x + \varepsilon. \end{array} \tag{3.12}$$

The system has a stable equilibrium point ($\dot{x} = \dot{y} = 0$) at $(x, y) = (0, 0)$ for $\mu < 0$ which becomes the centre of a stable limit cycle for $\mu > 0$. The Jacobian of this system is calculated:

$$J(x, y) = \begin{pmatrix} \mu(1 - x^2) & 1 \\ -1 & 0 \end{pmatrix}. \tag{3.13}$$

Evaluated at the stable point $(0, 0)$, the two complex-conjugate eigenvalues of the Jacobian are $\lambda = \frac{1}{2}(\mu \pm \sqrt{\mu^2 - 4})$. One therefore expects that the real part of both of the eigenvalues, equal to $\mu/2$, will approach zero as μ approaches zero from below, i.e. as the system approaches the bifurcation.

If integrated Equation (3.12) from $t = 0$ to $t = 400$ using a time dependent parameter

$$\mu(t) = -0.35 + 0.001t, \tag{3.14}$$

the bifurcation $\mu = 0$ will occur at $t = 350$. One can estimate the first eigenvalue of the Jacobian in a sliding window of 100 points. The mean (with error bars of one standard deviation) of 40 such experiments is shown in Figure 3.2. Note that the real part of the eigenvalue increases as expected and provides an early warning signal of the upcoming Hopf bifurcation.

3.2 Dimension reduction using EOF

The technique of dimension reduction can use empirical orthogonal functions. This is widely used in geosciences and meteorology to study large multivariate data sets. [Held and Kleinen, 2004] used the EOF technique prior to applying the ACF(1)-indicator to a set of gridded time series.

The EOF method involves projecting a data matrix Y of N time series onto a different basis to obtain N new time series. This basis is orthogonal and is constructed so that projecting Y onto the first basis vector maximises the variance of the projection. The second basis vector is such that projecting Y onto it will maximise the variance given that it must be orthogonal to the first, and so on. It may only need the first one or two EOFs to capture most of the variance of the whole data set of hundreds of time series.

Consider N observations $\mathbf{x}_1, \mathbf{x}_2, ..., \mathbf{x}_N$ of a P-dimensional system, collected in the $N \times P$ matrix X:

$$X = \begin{bmatrix} \mathbf{x}_1^\top \\ \vdots \\ \mathbf{x}_N^\top \end{bmatrix} = \begin{bmatrix} x_{11} & \dots & x_{1P} \\ \vdots & \ddots & \vdots \\ x_{N1} & \dots & x_{NP} \end{bmatrix}. \tag{3.15}$$

The method requires a mean-centred series. It is common to replace the series $[\mathbf{x}_n]_{n=1}^N$ with a new set of series $[\mathbf{y}_n]_{n=1}^N$, were the i^{th} element of $\mathbf{y}_n$, which is denoted by $\mathbf{y}_n^{(i)}$, is given by

$$\mathbf{y}_n^{(i)} = \mathbf{x}_n^{(i)} - \frac{1}{N} \sum_{j=1}^{N} \mathbf{x}_j^{(i)}, \tag{3.16}$$

so that a new matrix Y is created from the N time series $\mathbf{y}_1, \mathbf{y}_2, ..., \mathbf{y}_N$, similar to the matrix X in equation 3.15. Y is then projected onto a different basis to obtain the matrix T of EOF scores, the first column of T being the first EOF score. The matrix W of the basis vectors is made up of the eigenvectors of the covariance matrix C, where

$$C = \frac{1}{N-1} Y^\top Y. \tag{3.17}$$

The eigenvector corresponding to the largest eigenvalue is the first column of the matrix W, and so on. Then one can calculate $T = YW$. Often, only the first one or two EOF scores are required, therefore it is possible to create a smaller matrix W whose columns are the first one or two eigenvectors of C (those corresponding to the one or two largest eigenvalues).

Chapter 4

Applications

4.1 Artificial Data

4.1.1 Combination of EWS indicators

To illustrate how different indicators can be combined for analysis of tipping dynamics, in Figure 4.1 we show artificial data, in which a sigmoid function is superimposed with red noise with fluctuation exponent $\alpha = 0.7$, which was generated using a perturbation in the Fourier domain [Makse et al., 1996]. This dataset simulates a forced transition, where the pattern of fluctuations does not change but the underlying trend drives a system to another state. When applying the ACF-indicator without detrending, we see clear indication of critical behaviour, with values approaching 1. However, if we apply ACF-indicator with detrending or DFA-indicator, we can see that there is no real criticality in the data, and thus we can conclude that there is no genuine bifurcation or changing number of states in the system.

[Ditlevsen and Johnsen, 2010] suggested that autocorrelation alone is not enough to detect bifurcations, and in addition changes in variance should be monitored. Moreover, [Lennartz and Bunde, 2009] showed that in finite-size time series of length N with power-law correlation exponent γ, the variance $\Delta_N^2(s)$ and autocorrelation function $C_N(s)$ satisfy $\Delta_N^2(s) = C_N(s/s_x) + O(s^{-1})$, where $s_x = \left(\frac{2}{(2-\gamma)(1-\gamma)}\right)^{1/\gamma}$.

This confirms our observation that the DFA-indicator is meaningful in monitoring changes in the data alongside the ACF-indicator. For example, most climatic time series possess power-law correlations, and in the presence of memory, especially with high fluctuation exponent and nonstationarities, lag-1 autocorrelations are not informative. In Figure 4.2, we show data constructed of 11 chunks of red noise with increasing fluctuation exponent from 0.5 to 1.5, and one can see how the data becomes non-stationary, which is indicated by the trend in both indicators, as well as in the variance. Note that for the last two chunks, when the ACF-indicator reaches its maximal value 1, the DFA-indicator is able to detect further variability in the data, providing additional information about the changes in the variance.

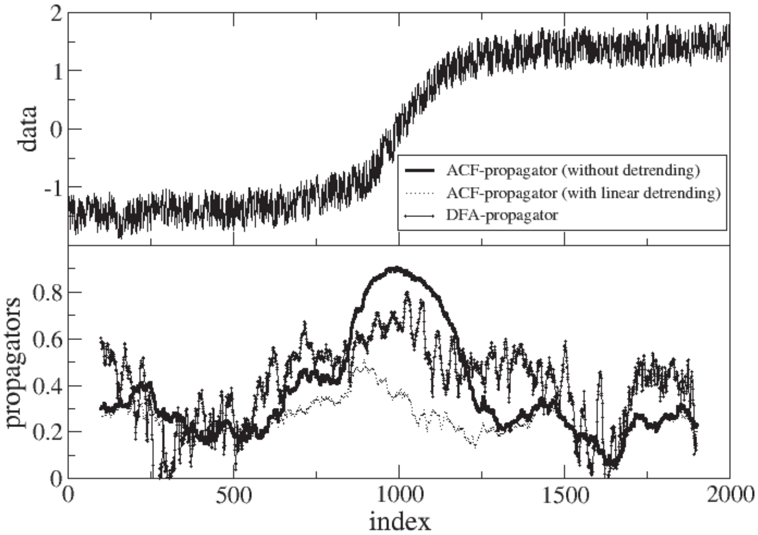

Figure 4.1: Sigmoid function superimposed with colour noise with fluctuation exponent $\alpha = 0.7$ (upper panel). ACF-indicators with and without detrending; the influence of the transitional trend is indicated in the area where these two curves of ACF-indicators differ. DFA-indicator, similarly to the ACF-indicator with detrending, shows no presence of a genuine bifurcation.

4.1.2 Non-potential system dynamics

Analytically, proving non-potentiality would require proving non-existence of non-trivial solutions to particular equations for the Lyapunov function (which is also difficult to find). However, potential analysis allows us to detect experimentally a sample without a globally existing potential, which is indicated by variable patches in the potential contour plot (see Sample 9, Figure 4.3). This intermittent pattern in the contour plot is caused by a varying potential or by complicated noise or both. In fact, series 9 may be non-potential with damping.

In the blind experiment [Livina et al., 2012] demonstrated in Figure 4.3, we show the indicators of the simulated data, and we note that, similarly to the test data, the indicators are able to detect the abrupt transitions. The main patterns of the simulated potential plots are the same as for the test data, and the non-potential character of the last sample is confirmed by the same patchy behaviour as in the test sample.

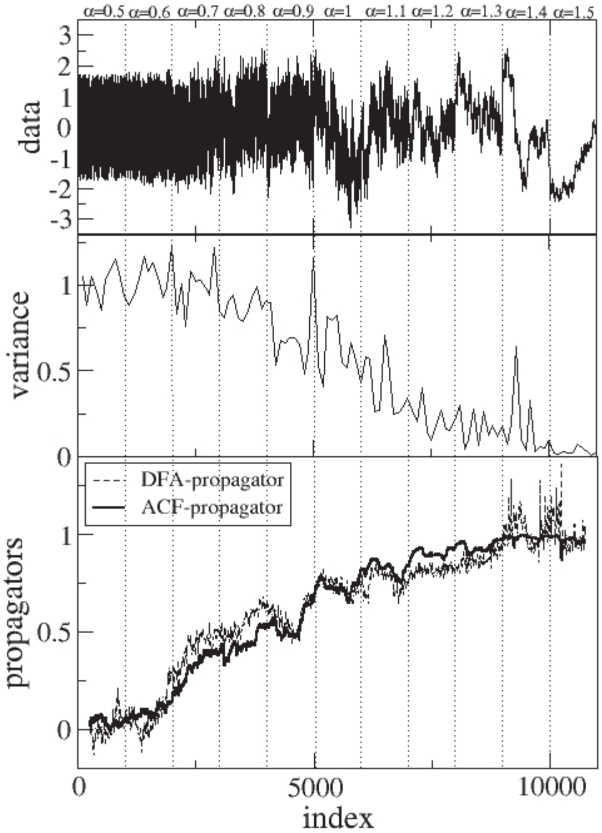

Figure 4.2: Artificial data combined of eleven chunks of 1,000 datapoints each with increasing fluctuation exponent α from 0.5 to 1.5 with step 0.1. Upper panel show data with increasing nonstationarities due to increasing memory. Middle panel shows variance calculated along the series in non-overlapping windows of size 100; due to the increasing memory the variance is decreasing. The bottom panel shows ACF- and DFA-indicators correctly detecting the critical behaviour of the data. Note that when ACF-indicator reaches critical value 1 and stabilises there, DFA-indicator is still able to detect variability in the data, and its fluctuations corresponds to changes in variance in the middle panel.

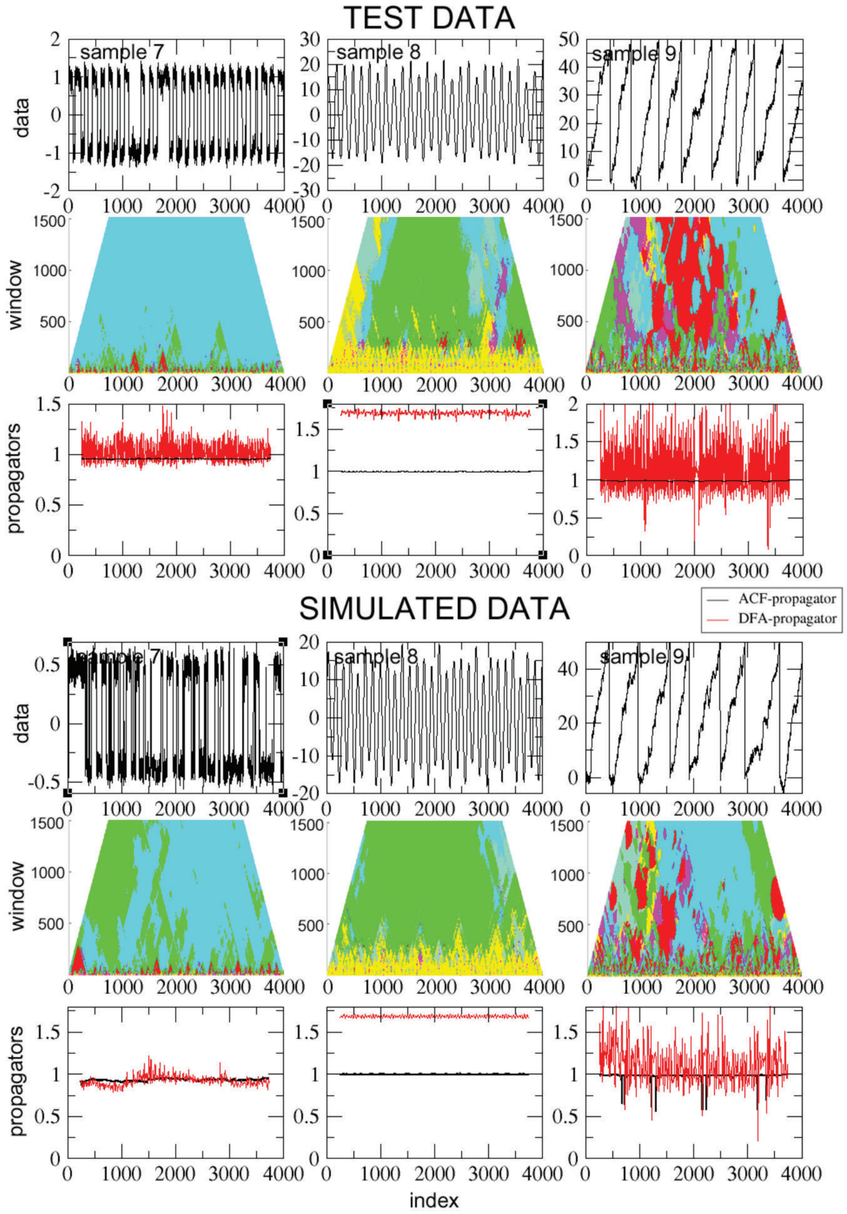

Figure 4.3: Three samples of artificial data in the blind test, and their potential colour plots and indicators. In the contour plots, detected one-well potential is denoted by red colour, double-well potential by green colour, triple-well potential by cyan colour, four-well-potential by purple colour; the number of detected wells bigger than 4 is denoted by yellow colour. ACF- and DFA-indicators were calculated with a fixed sliding window of size 500.

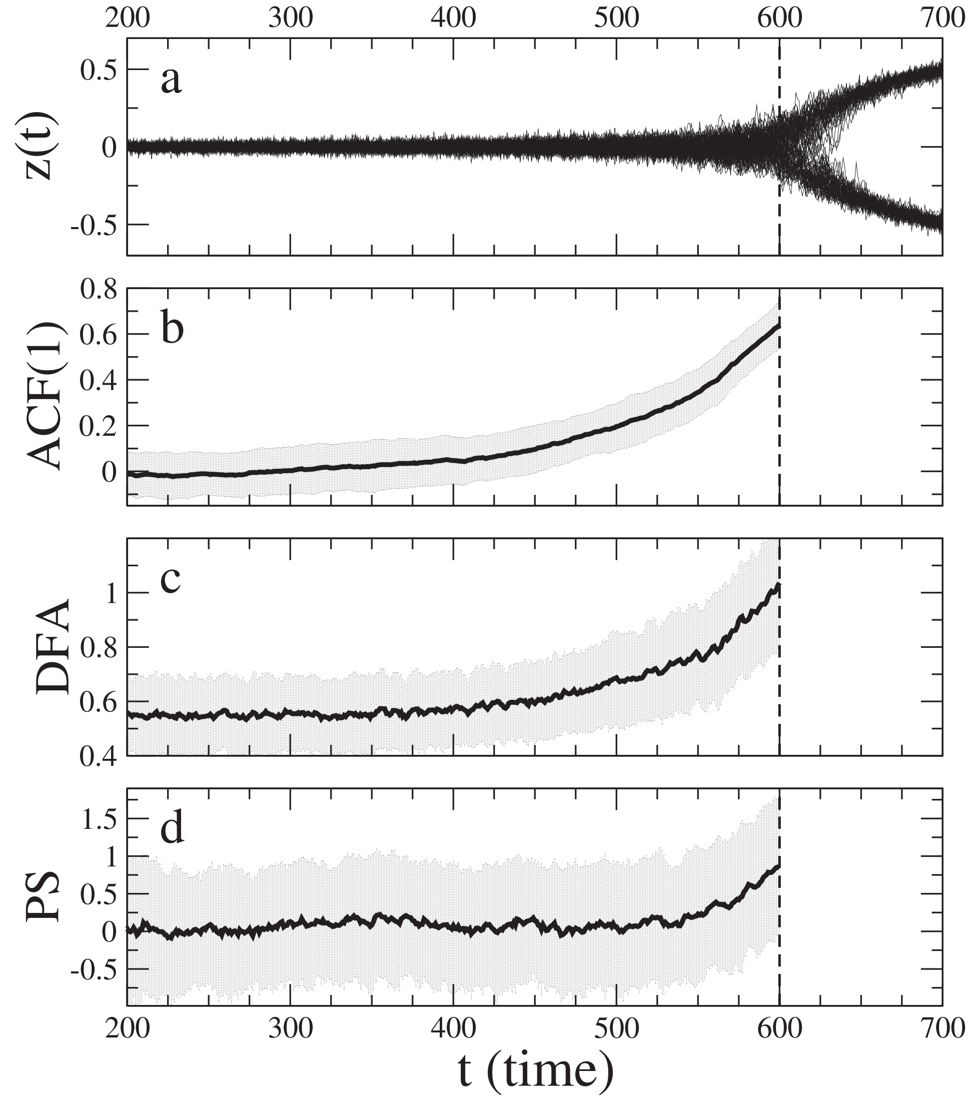

Figure 4.4: ACF(1) and PS-indicators applied to model data. (a) Data from 100 runs of the model (see eq. 4.1); (b,c,d) The mean ACF(1), DFA and PS-indicators (window size 100), with error bars of one standard deviation. The values of all the indicators begin to rise before the bifurcation, as expected, at around $t = 500$, although it is most noticeable for the ACF(1)-indicator.

By applying the PS-indicator to a system with a bifurcation point, one expects to observe behaviour comparable to the ACF(1)- and DFA-indicators. Consider a typical bifurcating system described by the stochastic differential equation

$$\dot{z}(t) = -\frac{\partial}{\partial z}\left(z^4 + \left(3 - \frac{t}{200}\right)z^2\right) + \sigma\eta_t, \tag{4.1}$$

where $\sigma = 0.05$ and the η_t are independent samples from a Gaussian distribution. Equation 4.1 is integrated numerically with a time step of δt = 0.05 from $t = 1$ to $t = 1000$, and sampled with $\Delta t = 1$. In the range $t < 600$ the system has a single stable node $z = 0$, the system then bifurcates after $t = 600$ into a double-well potential with stable nodes at $z = \pm 0.05\sqrt{t - 600}$, $z = 0$ becomes unstable.

The system is integrated 100 times to produce 100 time series of length 1000, see Figure 4.4a. The EWS indicators are then applied to the 100 time series and the mean of all 100 indicator series is plotted with error bars of one standard deviation. This can be performed with the ACF(1)- and DFA-indicators (Figure 4.4b,c) and the new PS-indicator (Figure 4.4d). The ACF(1)-indicator provides a clear EWS as expected, there is a definite increasing trend starting at around $t = 400$. Although the variance is larger, a similar trend with the PS-indicator is observed, so the validity of the PS-indicator as an EWS is illustrated.

4.1.3 Tipping forecasting: potential and ML-based

It is possible to fit models and train neural networks, such as the Long Short-Term Memory (LSTM) deep learning network, on historical data to forecast further dynamics. This is particularly important in projecting dynamics of time series that experience critical behaviour. However, transitions and bifurcations are difficult to forecast long-term, especially with complex non-stationarities.

For LSTM implementation, the Matlab deep neural network was used, with 90% of data used for training and 10% for testing. The following hyperparameter values were used: 200 hidden units, maximum number of epochs 250, initial learning rate 0.005, learn-rate-drop period of 60 epochs, and learn-rate-drop factor 0.2. The latter two parameters define the number of epochs that passes between adjustments to the learning rate during training and the factor by which the learning rate drops during training.

An alternative approach is to use the conventional time series models, such as the autoregressive integrated moving average (ARIMA) model. For ARIMA forecast, the Matlab implementation minimising Akaike and Bayesian information criteria was used to obtain the optimal model implementation based on the data input. In a similar manner, the technique of potential forecasting is data-driven and produces forecast time series using the statistical properties of the data, such as scaling dependencies and auto-correlations [Livina et al., 2013].

First, stationary Gaussian white noise was used with application of the three models to forecast the last part of a time series. This is not a bifurcating time series, and there is no building up of critical behaviour (see Figure 4.5), so it is an illustration of application of these models for later comparison with other two types of artificial data, where there are nonstationarities and a change of system states.

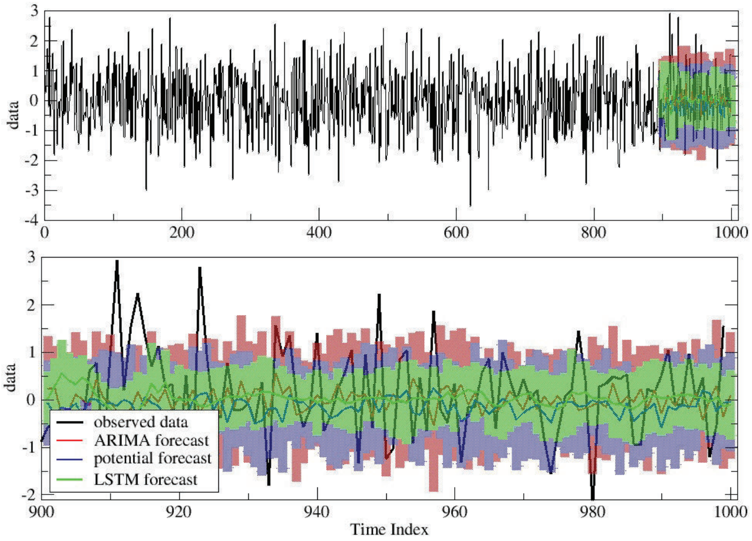

Figure 4.5: Three forecasts for Gaussian white noise. The time series has scaling exponent 0.5, and the modelled output is shown as averages and shaded areas of standard deviations over 20 trajectories for each model. It can be seen all three models produce trajectories with the same patterns as the considered time series.

The time series with power-law correlations with scaling exponent 1.1 was used. Figure 4.6 shows hindcasts of three techniques, ARIMA, potential forecasting and LSTM. With each technique, 20 stochastic trajectories were generated and used to obtain an ensemble forecast with an uncertainty range. The ensemble averages and their uncertainty ranges are shown by lines with shaded areas of different colours.

This non-stationary time series represents an example of possible transitional behaviour, and all three models reproduced the stationary pattern of the interval preceding the range of the forecast. The wide error range was the result of multiple trajectories with noise oscillations, and these were also quite short — these two factors defined the variability. This is one of the shortcomings of stochastic forecasts, and this should be considered in forecasting studies.

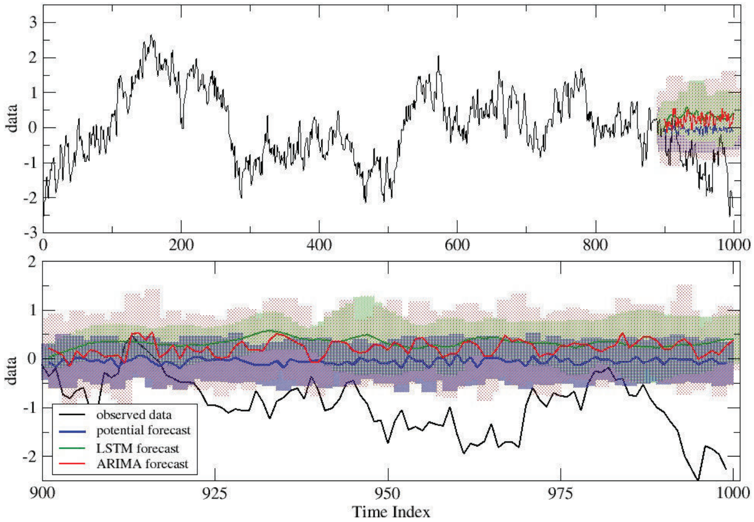

Figure 4.6: Three forecasts for the power-law correlated artificial data. The considered time series has scaling exponent 1.1, and modelled output is shown as averages and shaded areas of standard deviations over 20 trajectories for each model. ARIMA and potential forecasts were able to capture the fluctuations of the observed data, LSTM on the other hand was quite static.

The example was intentionally designed with such time series to illustrate the shortcomings of all forecasting techniques when dealing with highly auto-correlated data.

Nonstationary time series are notoriously difficult to forecast, but in the case of gradual potential dynamics, it is possible to obtain ensemble forecasts with satisfactory errors. Figure 4.7 shows double-well potential data with decreasing noise level leading to bifurcation (change from two to one potential well, as described in [Livina et al., 2011]).

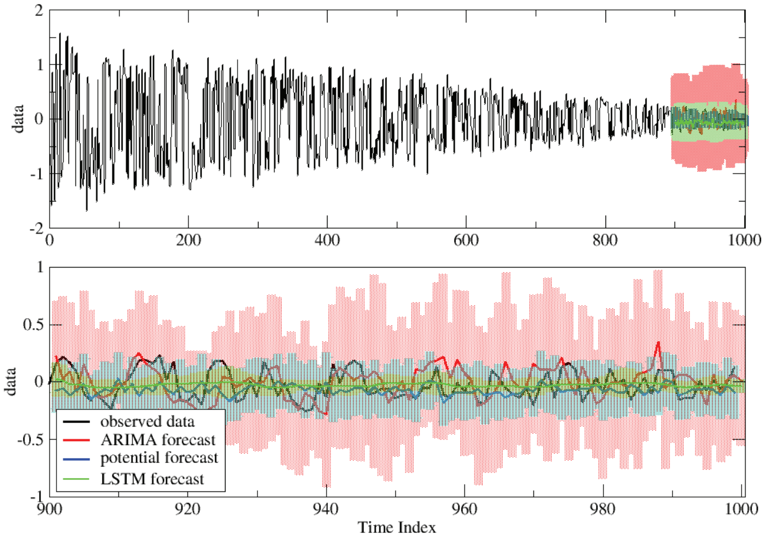

Figure 4.7: Three forecasts for the artificial double-well-potential data with decreasing noise level. Because of the gradual dynamics of the bifurcation, all three models perform satisfactory in thein ensemble forecasts.

4.2 Geophysical Data

The climate system is a nonlinear dynamical system which can possess multiple states, with shifts between these different climate states induced by stochastic forcing. The state variable z in Eq.(2.8) represents a large-scale climate variable, for instance, a paleo proxy (or observational aggregate) measure of temperature over a large region.

4.2.1 Oxygen isotope data

We examined ice-core proxy records of palaeotemperature at two different sites 325 km apart in Greenland; GRIP [Dansgaard et al., 1993] (Figure 4.8a) and NGRIP (Figure 4.8c), on the most recent GICC05 time scale [Svensson et al., 2008]. The data are δ^{18}O stable water isotope records, which are a proxy for past air temperature at the ice-core sites. The records can also be influenced by changing water source temperatures and snowfall seasonality. There are significant differences between the two records during the last ice age, with NGRIP systematically colder (depleted in ^{18}O), and more variable, when Northern Hemisphere ice sheets are more extensive. This may be because NGRIP received a greater fraction of colder air coming over the northern side of the Laurentide ice sheet. Importantly, the difference between the two records does not show a component of millennial variability. Here we consider the last 60 kyr NGRIP series with resolution 20 yr and a part of the GRIP series (which is 112 kyr long) in the same interval (60–0 kyr BP) with the same temporal resolution of 20 yr. For the detection of changing numbers of states we use sliding windows of varying length through each dataset.

To further test initial results obtained on the $\delta^{18}O$ data, we also examined Ca data from the GRIP project (Figure 4.9) that has annual resolution and spans the interval 91–11 kyr BP. In our analysis, we consider a subset of this dataset starting at 60 kyr BP to make it comparable with the other two datasets.

We visualize the estimated number of system states as a colour contour plot, expressed as a function of the time at the middle of the data window (x-axis, aligned with the data time scale) and the time length of the data window (y-axis). As the results are plotted at the middle of their corresponding time windows, each point in the contour plot should be compared with the segment of the time series data centered at that particular point, and having the length of the corresponding time window.

The GRIP (Figure 4.8a) and NGRIP (Figure 4.8c) palaeotemperature data are highly variable and non-stationary. Nevertheless, the algorithm detects a number of interesting common features in both records (Figure 4.8b,d). Over 60–25 kyr BP, two climate states are most commonly detected, consistent with the conventional view of the Dansgaard-Oeschger events. Exceptions of 1-state detection appear to be caused when long intervals of steady cooling, e.g. circa 55–50 kyr BP, occupy a significant part of the analysis window.

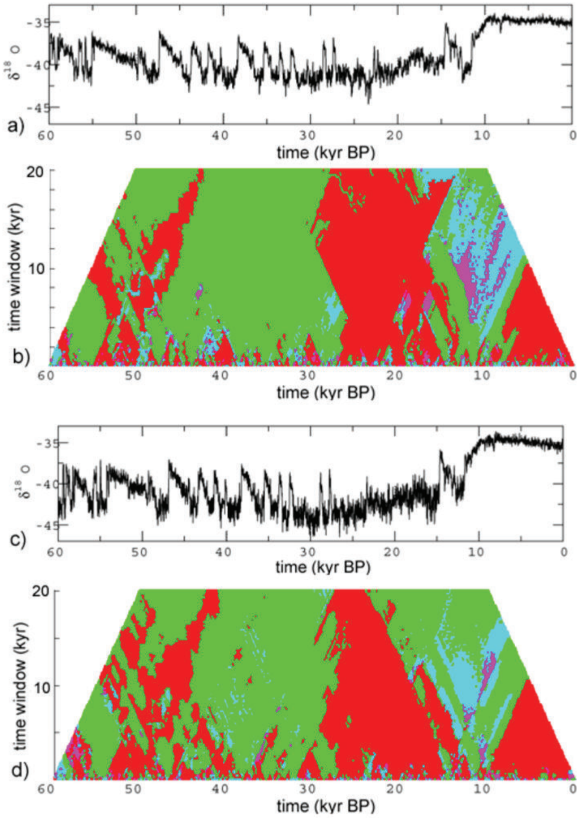

Figure 4.8: Potential analysis of Greenland ice-core data: **(a)** GRIP time series; **(b)** contour plot of the number of detected states versus time and size of sliding window, mapped at the end of the sliding time windows (red – one well, green – two, cyan – three, purple – four). **(c)** NGRIP time series; **(d)** contour plot of the number of detected states versus time and size of sliding window, mapped at the end of the sliding time windows, as in (b).

The most pronounced feature, common to both records, is a change from 2-well to 1-well potential, generally detected by 25 kyr BP and inferred to have occurred somewhat prior to its detection. The transition is most sharply detected in the GRIP data. In the NGRIP data, the transition is less pronounced, consistent with the greater noise level in this dataset obscuring the detection of states. In both records, the shift to a 1-well potential is persistent, indicating a bifurcation in the climate system that occurred late in the last ice age but prior to the Last Glacial Maximum (LGM) 23–19 kyr BP [Yokoyama et al., 2000].

The return of a second state in the climate is only detected in both records, for most window sizes, around 12 kyr BP (i.e. at the end of the Younger Dryas), although it is sporadically detected earlier, from the time of the Bolling warming, especially in the NGRIP data. This reflects the growing influences of changes during the deglaciation on the algorithm. The detection of up to four states as the data window spans the interval of the last glacial termination is probably an artefact of the nonstationarity in the time series. Although these four climate states plausibly represent the LGM, Bolling-Allerod warm interval, Younger Dryas cold interval, and the Holocene, it is perhaps more accurate to interpret the system as having two states about a moving trend.

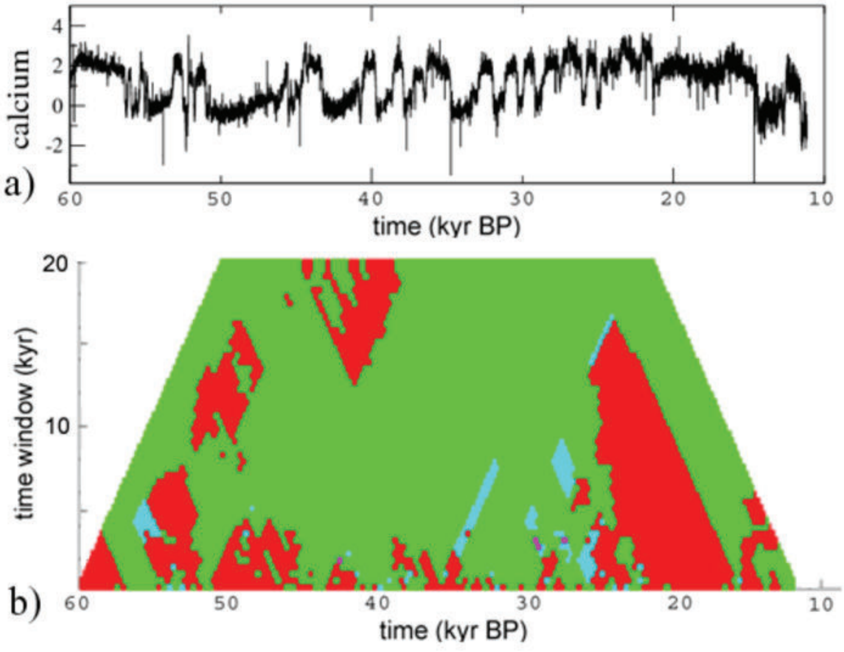

Figure 4.9: GRIP calcium data: a) time series; b) contour plot of the number of detected states versus time and size of sliding window, mapped at the middle of the sliding time windows (red – one, green – two, cyan – three, purple – four).

As the end of the analysis window advances into the Holocene, the estimated number of system states declines. However, the number of climate states only reduces to one for both datasets when the analysis window comprises only the data since around 10 kyr BP, i.e. the start of the Holocene. For the GRIP data, there is some indication that the 8.2 kyr BP event can influence the algorithm sufficiently that two climate states are detected when using small windows that span the event. The 8.2 kyr BP event is also visible in the NGRIP data but has poorer signal-to-noise ratio, so does not exert the same influence on the algorithm.

Given that the ice-core time series are much shorter (in terms of the number of recorded transitions between states) than the simulated time series used above, the question arises how robust our results on the ice-core data are statistically. The mean recurrence time (Kramers waiting time) between DO events is 2.8 kyr which is about the same as in the artificial data in nondimensional system time units. Hence the maximum window size of 20 kyr in Figure 4.8b,d corresponds to 20 units on the vertical axis in Figure 2.8c. This is just enough to enable a correct detection of the number of states with very high probability. Together with the fact that our results are robust for a range of window sizes between 10 and 20 kyr this gives us some confidence that our conclusions are reliable. Results in Figure 2.9 confirm the same using different techniques.

4.2.2 Arctic sea-ice data

Arctic sea-ice has experienced striking reductions in areal coverage [Nghiem, 2007, Stroeve et al., 2007], especially in 2007-2012, having the six lowest ice cover minima in the satellite record during summer (Figure 4.10). Observations have fallen below IPCC model projections [Stroeve et al., 2007], despite the models having been in agreement with the observations in the 1970s. The latest models are more consistent with satellite observations (1979-present), but still fail to capture the full extent of the observed downward trend [Stroeve et al., 2012]. Summer ice cover is forecast to disappear later this century [Boe et al., 2009], but the nature of the underlying transition is debated [Amstrup, 2010, Eisenman and Wettlaufer, 2009, Lenton et al., 2008, Lindsay and Zhang, 2005, Tietsche et al., 2011, Winton, 2006].

Arctic sea-ice has been identified as a potential tipping element in the Earth's climate system [Lenton et al., 2008], and at least one study suggests it has already passed a 'tipping point' [Lindsay and Zhang, 2005]. In the future, some models forecast abrupt ice loss events [Amstrup, 2010], on the way to a seasonally ice-free Arctic. These may qualify as passing tipping points following the broad definition given in [Lenton et al., 2008] of a point at which a small change in forcing leads to a qualitative change in the future state of a system. The definition includes both reversible and irreversible transitions, bifurcations and some non-bifurcation phenomena.

However, most recent papers on the Arctic sea-ice opt for a narrower definition of a tipping point, addressing whether summer sea-ice loss will involve an irreversible (e.g. saddle-node/fold) bifurcation. They find instead that in models the loss of summer sea-ice cover is highly reversible [Amstrup, 2010, Eisenman and Wettlaufer, 2009, Tietsche et al., 2011, Winton, 2006]. Abrupt ice loss events are then attributed to the loss of year-round sea-ice in the Arctic making the remaining ice more vulnerable to summer melt, and prone to larger fluctuations in area coverage [Notz, 2009]. An exception is a recent model [Abbot et al., 2011] showing that positive feedbacks involving clouds can create multiple stable states for seasonal ice cover and bifurcations between them. Furthermore, models of past abrupt climate changes in the Arctic have shown multiple stable states for sea-ice cover in the Barents and Kara Seas region and abrupt switches between them [Bengtsson et al., 2004, Semenov et al., 2009]. This suggests that sub-Arctic scale ‘tipping points’ in sea-ice cover are conceivable.

On viewing the satellite-derived daily record of sea-ice area from 1979 to present (Figure 4.10a), it is clear that the last five years have been characterized by an increase in the amplitude of seasonal sea-ice variation (Figure 4.10b). The annual ice cover minimum dropped by the order of $\sim 10^6$ km^2 more than the annual maximum in 2007 and the difference has been maintained since then (Figure 4.10c). This already suggests an abrupt and persistent change in sea-ice dynamics. It led to a hypothesis that the sea-ice may have passed a bifurcation-type tipping point, in which a new attractor for lower summer-autumn sea-ice cover became stable and began to be sampled in summer 2007, and in every summer since, with seasonal switches to/from the pre-existing attractor.

The sea-ice extent time series was derived by [Eisenman, 2010] on the basis of sea ice concentration using the NASA Team algorithm from Nimbus-7 SMMR (1978-1987), DMSP SSM/I (1987-2009), and DMSP SSMIS (2008-present) satellite passive microwave radiances on a 25km x 25km polar stereographic grid [Cavalieri et al., 1996, Maslanik and Stroeve, 1999, Meier et al., 2021]. During periods of instrumental transitions, the overlapping datasets were averaged. The extent was calculated by summing the areas of all grid boxes with at least 15% ice concentration. Details of the spatial data interpolation are given by [Eisenman, 2010]. The time series spans 1979-2009, and where it has 2-day resolution (when SMMR operated every other day in three months during the record, in 10/1978, 12/1987 and 1/1988), we interpolate to daily resolution to obtain a homogeneous time-series.

For both datasets – area and extent – the mean seasonal cycle over the first 30 years of data (1979-2008) was removed. We also examined the effect of constructing and removing a different averaging interval (1979-2011), which produces a very similar residual series, just vertically shifted along the y-axis, i.e. the dynamics of the residual fluctuations remained the same. Hence this gives similar results and we do not show it here.

After removing the mean seasonal cycle (1979-2008), the remaining fluctuations in sea-ice area include some of order 10^6 km^2 (Figure 4.11a). The largest anomalies are in 1996 (maximum of the series) and 2007-2011 (minima). They typically occur in the summer-autumn, when the sea-ice area is at its lowest in the seasonal cycle. Given the size of sea-ice fluctuations during 2007-2011 (Figure 4.11a) and the pronounced drop in sea-ice minima relative to sea-ice maxima since 2007 (Figure 4.10c), we considered whether the residuals exhibited an abrupt change to multi-modality in 2007.

On analyzing the residual sea-ice area fluctuations using our method of potential analysis, over long time windows (here $>$1 year), we typically find a single mode and corresponding attractor, representing the normal seasonal cycle of sea-ice variability (Figure 4.11b). Sometimes a second mode is detected associated with e.g., the sea-ice maximum in 1996, but these changes are not found simultaneously and persistently across a wide range of window lengths. However, from 2007 onwards, a persistent switch to two modes or attractors is detected, across a wide range of window lengths up to $>$10 years (Figure 4.11b).

The stability of the attractor(s) for the residual sea-ice fluctuations can be reconstructed, in the form of potential curves for fixed intervals of the data (Figure 4.11c), with associated error estimates (on the coefficients of the polynomial function describing the potential [Livina et al., 2011]). The sea-ice residuals are typically characterized by a single mode and corresponding attractor. The interval 1996-9 (including the 1996 maximum anomaly) shows signs of a second higher ice cover attractor that is degenerate (i.e. not fully stable). In 2000-3 there is a return to a single attractor. In 2004-7, which includes the record September 2007 sea-ice retreat, a low ice-cover attractor starts to appear in the fluctuations. Then in 2008-11 the potential separates into two attractors, although the error range allows for one or the other of these to be degenerate.

The potential curves are derived from histograms of the original data [Livina et al., 2011] (Figure 4.11d), which confirm a second mode appearing among a long tail of negative fluctuations during 2004-2007, followed by a separation of multiple modes during 2008-2011, which the method fits as a bi-modal distribution. Thus, we originally hypothesized that the Arctic sea-ice recently passed a bifurcation point [Livina and Lenton, 2013], which created a new lower ice cover attractor for the residual deseasonalised fluctuations. Since then it has fluctuated between its normal attractor for seasonal variability and the new, lower ice cover attractor.

However, an abrupt change in the amplitude of the seasonal cycle will leave a residual record that has some seasonality on one side of the transition or the other. These remnant seasonal fluctuations will in turn produce a bi-modal distribution, which is accurately detected by our method – hence care is needed over how to interpret this.

Prior to 2007 there is no consistent early warning signal of destabilization (Figure 4.12c,e,g). The indicators all increased around the anomalous sea-ice maximum in 1996, but then they all declined toward 2007, consistent with our potential reconstruction (Figure 4.11c). There was no consistent early warning signal of critical slowing down before the hypothesized bifurcation. Instead the sea-ice showed signs of increasing stability in the preceding decade, contrary to what would be expected from an approach to bifurcation.

The sea-ice retreat in 2007 caused abrupt increases in all the indicators, which have continued to rise since then (Figure 4.12c,e,g). Sensitivity analysis reveals a robust upward trend in the DFA-indicator across the whole dataset (Figure 4.12d), but no robust overall trend in the ACF-indicator or variance (Figure 4.12b,f).

An abrupt and persistent change in sea-ice dynamics is detected to have occurred in 2007. This involves an extra $\sim 10^6$ km^2 or more sea-ice loss each summer-autumn since then. Our initial hypothesis that this abrupt increase in the amplitude of the seasonal cycle of sea-ice variability occurred through a bifurcation mechanism [Livina and Lenton, 2013], is not consistently supported. Thus, the underlying causal mechanism remains uncertain. Still, there must be some amplifying positive feedback mechanisms contributing to the abrupt increase in summer-autumn ice loss.

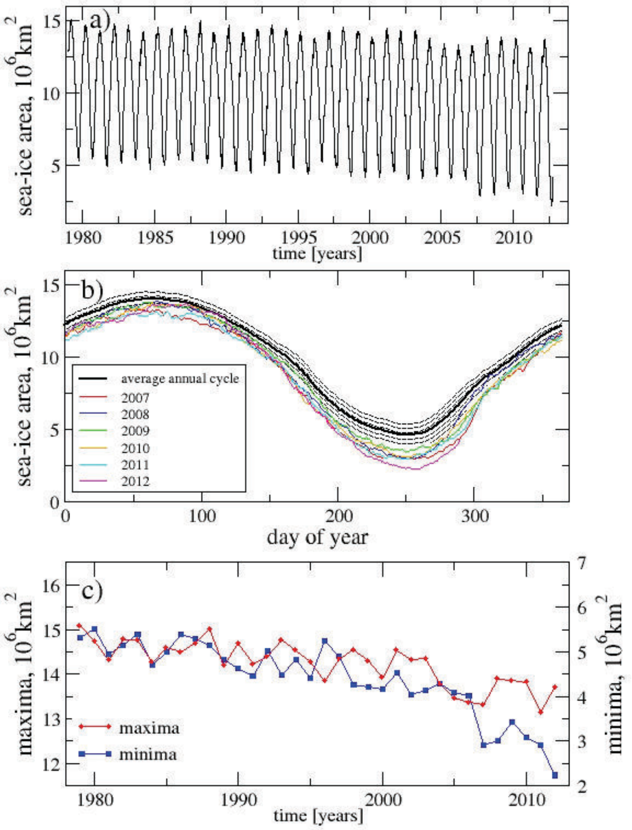

Figure 4.10: Arctic sea-ice area from satellite data. (a) Arctic sea-ice area, 1979-2012. (b) The mean annual cycle of the area data over 1979-2008 inclusive (solid line, shaded area denotes two error bars), together with the last five anomalous years. (c) Annual maxima (left axis) and minima (right axis) showing an abrupt increase in amplitude of the seasonal cycle in 2007.

The abrupt increase in the seasonal cycle that we detect clearly does not involve total seasonal sea-ice loss and hence is sub-Arctic in scale. However, there may be a precedent for this; past abrupt Arctic cooling and warming events have been linked to switches between alternative states for sea-ice cover in the Barents and Kara Seas region [Bengtsson et al., 2004, Semenov et al., 2009]. Such sub-Arctic-scale switches can still have significant impacts, indeed recent ice loss from the Barents and Kara Seas has been linked to cold winter extremes over Eurasia [Petoukhov and Semenov, 2010]. The connection between surface temperature, sea level pressure and winds in the Arctic region, and their effect on the sea-ice cover is discussed by [Comiso, 2012].

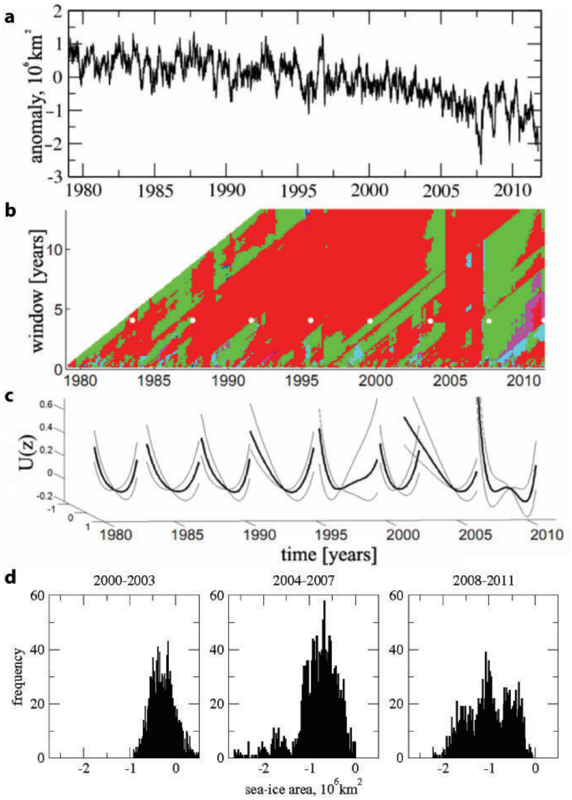

Figure 4.11: Analysis of Arctic sea-ice area. (a) Sea-ice area anomaly, daily data with mean seasonal cycle removed. (b) Contour plot of number of detected states, where red = 1 detected state, green = 2, cyan = 3, magenta = 4. Results plotted as a function of sliding window length at the end of the window. (c) Reconstructed potential curves of eight 4-year time intervals, corresponding to the white dots in (b). Here 'z' is sea-ice area fluctuation on a shifted scale. Faint lines are potential curves derived from error estimates on the coefficients of the polynomial potential function. In the penultimate interval 2004-2007 a second state starts to appear and in the final interval 2008-2011 there are two states of comparable stability. (d) Histograms of the data for 2000-2003, 2004-2007, 2008-2011 from which the corresponding potential curves are derived.

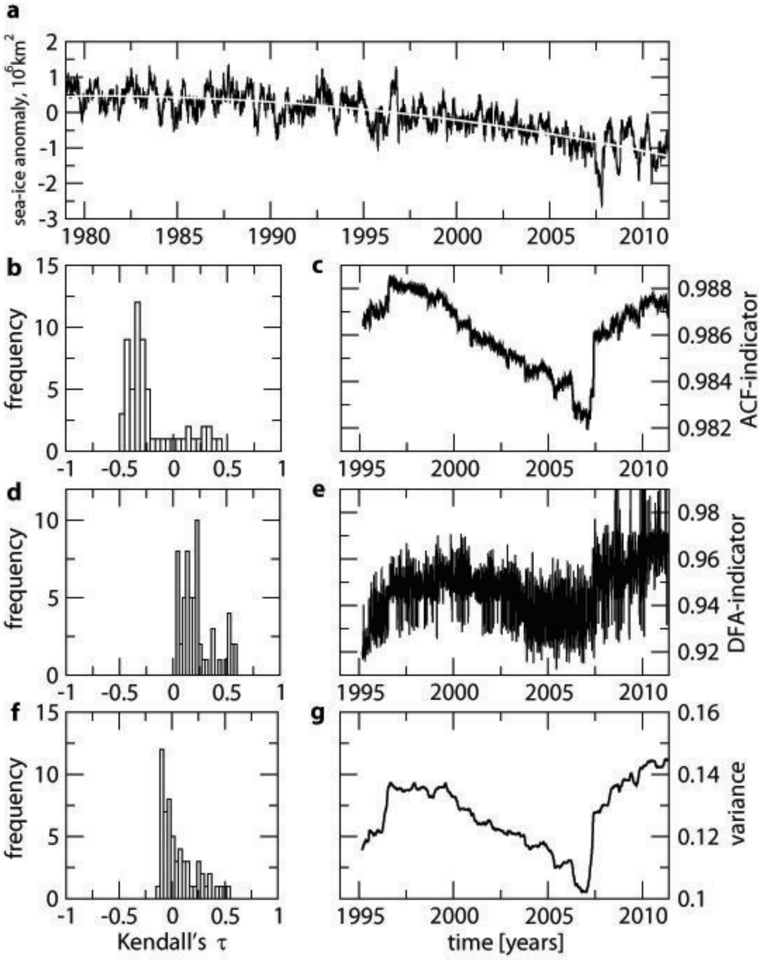

Figure 4.12: Search for early warning signals of bifurcation in Arctic sea-ice area data. (a) Sea-ice area anomaly showing the quadratic downward trend that is removed prior to calculating the instability indicators. Right panels show example indicators using a sliding window of length half the series, with results plotted at the end of the sliding window. Indicators from: (c) autocorrelation function (ACF), (e) detrended fluctuation analysis (DFA) and (g) variance. Left panels show histograms of the Kendall statistic for the trend in the indicators when varying the sliding window length from 1/4 to 3/4 of the series: (b) ACF-indicator, (d) DFA-indicator, (f) variance.

4.2.3 Ocean acoustic noise data

The Preparatory Commission of the Comprehensive Nuclear-Test-Ban Treaty Organization (CTBTO) has established a global network of underwater hydrophones as a part of its hydroacoustic observations (others being seismic, infrasound, and radionuclide), with the goal of continuous monitoring for possible nuclear explosions. The CTBTO database provides several unique and large oceanic acoustic records, covering more than ten years of continuous recording with a high temporal resolution of 250 Hz. In this section, we study the records of the hydrophone H01W1 at the Cape Leeuwin station. The hydrophone is located at a depth of about 1 km off the southwest shore of Australia.

We study the large CTBTO record (2003-2016) of the Cape Leeuwin hydrophone, series H01W1, which is a 250 Hz-sampled time series of ocean sound pressure. The raw data represents 3 Tb of binary waveforms, which after extraction constitute 95 billion points in the time series. We analyse 1-minute averages of sound pressure level (SPL) in five frequency bands (broadband, 10-30 Hz, 40-60 Hz, 56-70 Hz, and 85-105 Hz), of about 7 million points per time series. This data has pronounced seasonality and some small gaps, and therefore we perform interpolation and deseasonalisation of all five time series, the result of which can be seen in Figure 4.22.

The data samples were scaled using their calibration factors (provided by CTBTO), and an inverse filter of the recording system's frequency response was applied to eliminate the effect of the acquisition chain on the frequency response of the recordings. The Fast Fourier Transform of the signal was computed using rectangular windows of 15 thousand samples (i.e. 1-minute-long intervals at 250 Hz sampling rate) and the broadband signal was then filtered in 5 frequency bands (5-115 Hz, 10-30 Hz, 40-60 Hz, 56-70 Hz, 85-105 Hz) via selection of the corresponding FFT bins within each frequency band. Then the resulting sound pressure level in dB re 1 μ Pa^2 for each frequency band was calculated [ISO18405:2017, 2017, Robinson et al., 2014]. Finally, outliers, i.e. levels greater than 20 dB from the average of the entire time series of SPL values, were removed.

Because of the data gaps, we interpolate the SPL data to achieve equidistant 1-minute temporal resolution. We then remove the seasonal periodicity by subtracting the averaged seasonal cycle over the 14 years of observation to obtain the fluctuations

$$z_i = S_i - \bar{S}_i, \tag{4.2}$$

where S_i is the interpolated SPL data, and $\bar{S}_i$ is the mean 1-minute interpolated SPL data. The resulting fluctuations are shown in the right column of Figure 4.13, for the broadband and four selected sub-bands.

We next apply the pre-tipping analysis (early warning signals) to analyse lag-1 autocorrelations and variance of the broadband SPL record, with estimation of uncertainty. We vary the length of the sliding windows for calculating these indicators between 1/4 and 3/4 of the record length to obtain the averaged curves and standard uncertainties and display the indicator values at the end of each window, as shown in Figure 4.15.

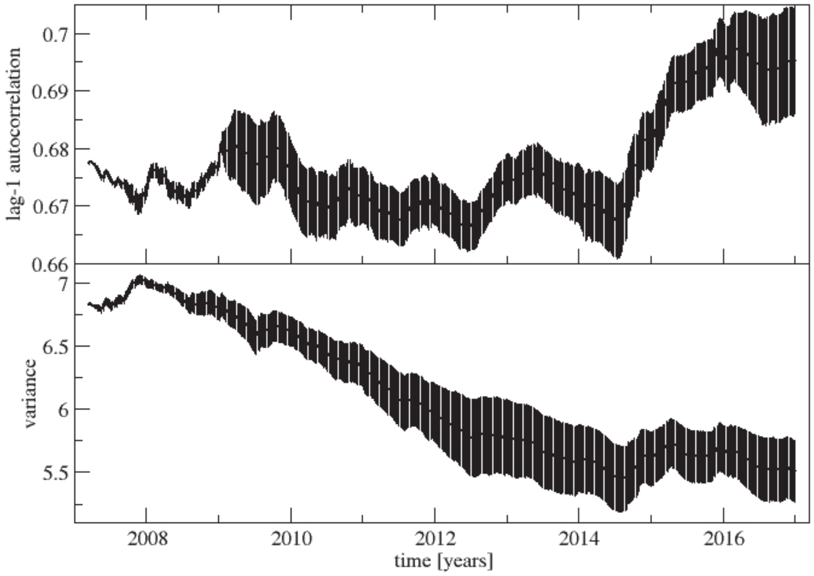

Figure 4.14: Early-warning indicators of the broadband SPL dataset: lag-1 autocorrelation (upper panel) and variance (lower panel), calculated with variable window lengths, from 1/4 to 3/4 of the record length, and corresponding standard uncertainties, displayed at the end of each window. Both indicators demonstrate nonstationary behaviour (increasing autocorrelation and decreasing variance), which denotes long-term development of a possible tipping in future.

The noticeable change at the end of these early-warning indicators may be related to the unusually large El Niño event of 2015-2016. One can see that the variance decline slows down and autocorrelation sharply rises, which means that the increase in memory is not accompanied by increasing amplitude of acoustic fluctuations. Such effects may happen when a dynamical system experiences critical slowing down prior to a bifurcational tipping. As we hypothesise that the El Niño signature may be related to changes in both oceanic dynamics and fauna, the increasing memory in the acoustic data may reflect, for instance, the observation that during the El Niño the Cape Leeuwin current slows down [Feng et al., 2003]. The slower ocean current introduces more inertia in the dynamical system, and therefore higher temporal memory/autocorrelations.

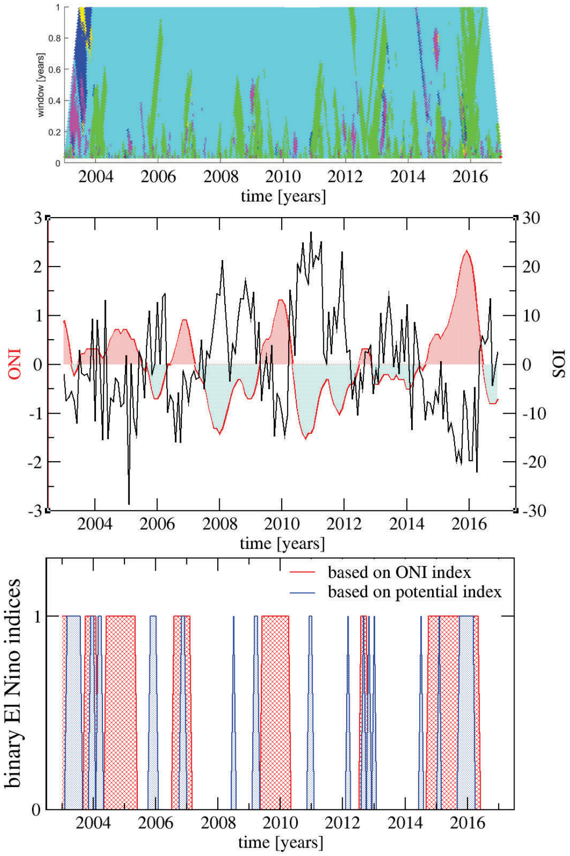

Figure 4.15: Upper panel: potential analysis plot of the broadband SPL data, with varying window length (y-axis) from one day to one year. The colours denote the number of detected potential states: green – two, cyan – three. Specks of magenta denote very short periods of a higher number of states, which correspond to highly variable (possibly non-potential) subsets of data of small size. Middle panel: ENSO indices ONI and SOI, known to be anti-correlated, which indicate several ENSO events (El Niño and La Niña). These can also be seen in the potential analysis plot. ONI positive and negative values (roughly corresponding to El Niño and La Niña) are shaded by light red and light blue respectively, for better comparison with the indices in the lower panel. Lower panel: binary indices derived from the ONI index and potential plot (at the level 0.5 year of y-scale in the top panel) which have values 1 when there is an El Niño (in the case of ONI-based binary index) or anomalous potential state (in the case of the potential binary index).

Similarly to the CTBTO data, the effect of increasing autocorrelation and decreasing variance was earlier observed in bifurcating artificial data changing from white noise to random walk, in [Livina et al., 2012]. The acoustics dynamics may be undergoing a similar tipping. Note that this analysis of early warning signals is performed with large enough windows (starting from length of 3 years up to 9 years), which identify large-scale variability, with possible dynamics on the scale of decades ahead.

Further, we apply potential analysis to identify smaller-scale variability, varying the length of the sliding window from three days to one year. The resulting potential plot is shown in Figure 4.15.

El Niño/Southern Oscillation (ENSO) can be monitored using several indices, which are obtained by averaging climatic variables to make the presence of El Niño more visible in the series. We show in Figure 4.15 two of them: Southern Oscillation Index (SOI) and Oceanic Niño Index (ONI). SOI is based on the sea level pressure differences between Tahiti and Darwin, Australia. ONI is based on the 3-month running mean of sea-surface temperature anomalies ERSST.v4 SST [Huang, 2014] in the Niño 3.4 region (NOAA SOI index). Negative SOI (positive ONI) correspond to El Niño events, characterised by warm SST in the eastern and central tropical Pacific [Trenberth and Caron, 2000].

We calculate, for easier comparison of El Niño indices and potential analysis, two binary indices derived from the ONI index and from the single level of the potential plot at the scale 0.5 yr. The bars in the bottom panel of Figure 4.15 show the occurence of El Niño events in the ONI index (which is less noisy than SOI), and at the same time we plot a binary index showing periods when the system potential does not follow its “regular” three-well-potential pattern; these two indices have agreement in several periods corresponding to the known El Niño events (2003, 2004, 2006, 2009, 2015-2016), which illustrates our hypothesis of the El Niño signature in the acoustic data.

The vertical span of the features of the potential plot (the specks of different colours) correspond to the time scale of the change, i.e. the size of the time window, within which the change has been detected. As El Niño is a seasonal phenomenon (except the unusually long event of 2015-2016), most of such specks are located within the window of size one year. The large event of 2015-2016, indeed, extends higher than that. To address this time scale, we derived the binary potential index using the detection data at fixed time scale of 0.5 year, at which most El Niño events should be present in the detection statistics.

We do not claim that the potential colour plot could be used for early warning signals (such as prognosis of El Niño), in this system or in others. Moreover, there may be other factors causing structural changes in the acoustic data, rather than El Niño or La Niña. On the other hand, detection of such changes, indeed, can be useful for other studies that could investigate attribution of structural variability, and here the technique of potential analysis might be very useful.

It is known [Feng et al., 2003] that the Leeuwin Current is influenced by El Niño, which causes lower temperature and slower current. This causes a number of climatic and environmental changes (including the impact of El Niño on sea level, current transport, and migration of marine animals), and this may affect the acoustic signal. In particular, the local sea bottom slopes near Cape Leeuwin are very steep, with large underwater peaks, which may be inducing reflection and scattering of the acoustic signal at greater depths, where the hydrophone is located.

Because the considered data is obtained from a hydrophone oriented in the west direction (towards the Indian ocean), one would expect very little influence on distant deep sources in predominantly west or south-westerly direction as their ray paths sample mostly the non-shallow ocean far removed from the coast and only passing through the Leeuwin Current at the last stage. But for surface-originating sounds from other directions there may well be an impact: where the sea bottom slopes through the SOFAR channel, particularly if it does so steeply, surface sounds and seismic waves can be reflected or refracted into the SOFAR channel. Note that on account of the lower speed of sound in water compared to rock, refraction is towards the normal for seismic waves coupling into water so such coupling is not efficient for a mostly horizontal sea bottom: the hydrophone array will predominantly see seismic waves that impinge close-to-vertically from below (have a small slowness, high apparent velocity across the array) which are subsequently scattered by the wavy sea surface instead of propagating coherently onwards. Hence steep slopes couple better.

The current dynamics of the acoustic fluctuations, which demonstrate slow but steady changes in early warning indicators, gives indication of an upcoming tipping point in this hydroacoustic system, with possible appearance / disappearance of system states, which in this context denote higher / lower SPL fluctuations.

4.2.4 Relative humidity data

Dew point and humidity variables are important for monitoring climate. Unlike, for instance, temperature, they require complex sensitive instruments and data processing for obtaining measurements [Willett et al., 2014]. Analysis of these variables may provide rich insight in both instrumental and climate dynamics. In this section, we apply tipping point analysis to obtain various signatures of the dynamics of the climate system and report the obtained results.

Using the dewpoint variable and pressure variable from the ECMWF reanalysis data, we obtained the relative humidity variable from measured dewpoint following [Willett et al., 2014]. The data were pre-processed as follows:

- Where large gaps were observed, the initial part of the data that included the gap was removed,
- Global trend and seasonality were removed using singular spectrum analysis (SSA) from both dewpoint and relative humidity variables,
- Detrended fluctuations of dewpoint and relative humidity were analysed using lag-1 autocorrelation for early warning signals, with uncertainty estimates based on varying size of the sliding windows: 0.3 to 0.7 of the time series length. The purpose of this analysis is to identify global warming trend in the data. The strength of trend was quantified using Mann-Kendall coefficient for assessment of monotonous trend. For the considered 55 stations, averaged Kendall values of dewpoint detrended fluctuations are 0.30 ± 0.54, whereas for relative humidity detrended fluctuations, the Kendall values are 0.68 ± 0.44.

Detrended fluctuations were also analysed with a short single window of 0.1 of each time series length. The purpose of this analysis, with higher sensitivity of the EWS indicator in smaller windows, was to identify nonstationarities in the data, possibly caused by instrumental artefacts, reversed short-term transitions and global trends.

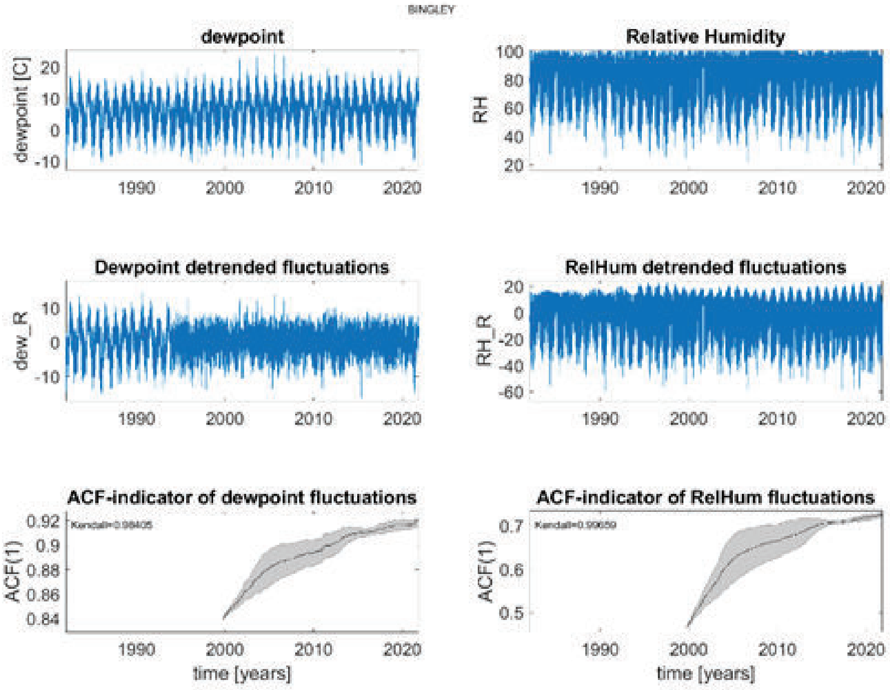

Figure 4.16: Dewpoint and relative humidity data, with long-window EWS indicators.

Figure 4.16 illustrates the analysis of global warming trends, with uncertainty quantification (shaded area) and Mann-Kendall trend coefficient [Lenton et al., 2012, Livina et al., 2014].

The same ACF1-indicator can be applied with a short sliding window, which allows one analyse short-scale transitions, such as those caused by instrumental changes. For example, in the 1980s-1990s there were replacements of the earlier analogue instruments with modern digital, which can be detected using ACF-based indicators, as can be seen in Figure 4.17

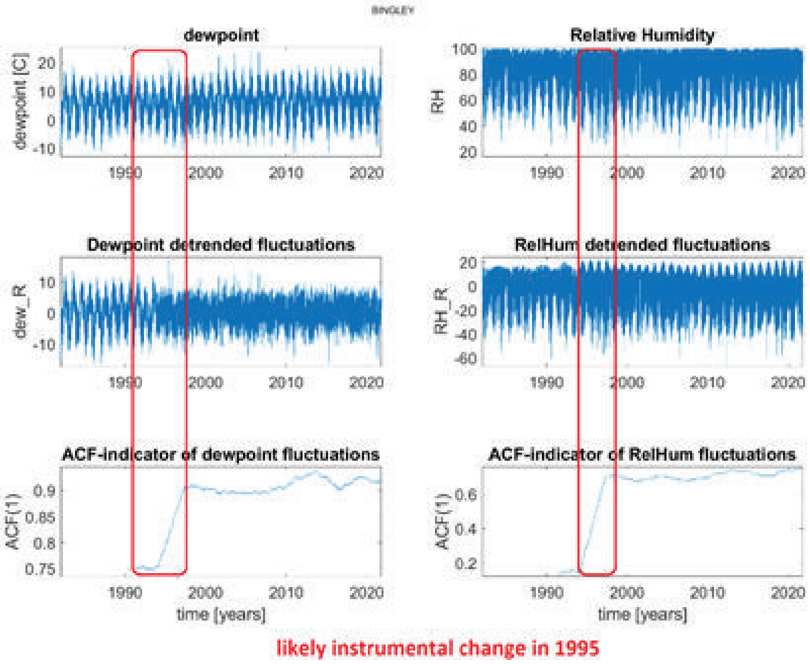

Figure 4.17: Dewpoint and relative humidity data, with short-window EWS indicators.

One of the goals of using EWS indicators in this study is automatic identification of critical transitions due to instrumental changes. This can be done using probabilistic estimation of change points in the EWS indicators. By applying the Matlab package ”Bayesian Changepoint Detection & Time Series Decomposition”, we can identify change points in EWS indicators, specifically those of statistical significance, which will allow automatic detection of such changes. As demonstrated in Figure 4.18, the changes in the indicators that are statistically significant can be identified in the probabilities of the changes based on the Bayesian ensembles with Monte Carlo simulations[Zhao et al., 2019]. The considered variable was detrended relative humidity, for which the ACF-indicator was calculated, which then was analysed using the Bayesian model ensemble. The changes with statistically significant peaks $(p > 0.8)$ are narrow, and the change events can be timed with small uncertainty of several weeks, if necessary.

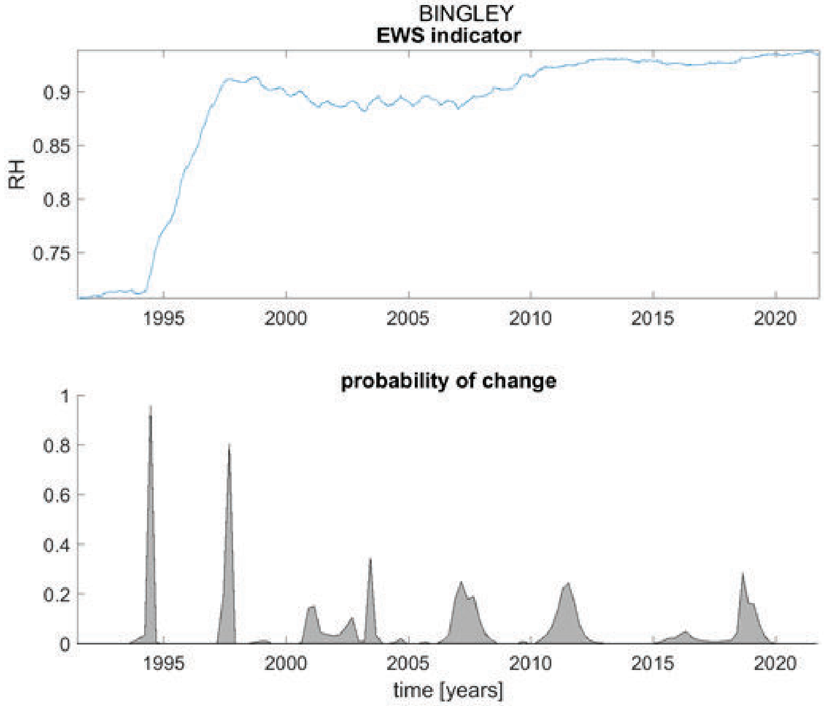

Figure 4.18: Upper panel: EWS indicator of relative humidity and detection of its change points. Lower panel: statistically significant changes are denoted by high probabilities in Bayesian ensemble modelling.

In the case of the weather station at Bingley, there are two statistically significant changes, in 1994 and in 1998, which is indicated in the lower panel with probabilities of changes.

4.3 Industrial data

4.3.1 NPL footbridge data

The footbridge at National Physical Laboratory (Teddington, UK) was a reinforced concrete structure built in the 1960s. It was used in a recent NPL research project to create a full-scale test bed, designed to investigate the capability of sensing systems for long-term monitoring in outdoor conditions.

The main focus of the project was on sensors and techniques suitable for the assessment of the performance of reinforced concrete structures under various loading conditions in outdoor environments during accelerated tests. The collected data formed the test set for examining various numerical techniques.

The unique aspect of this project is that a typical 1960's structure, aged naturally while in use for nearly 50 years, was then subjected to series of damage/repair cycles with progressively increasing levels of damage. There were two main research goals. The first one was to monitor the performance of various sensing technologies and their responses to different levels of damage applied to the bridge. The second was to use a selection of advanced numerical methods of data analysis and explore their capability for damage detection at relatively early stages.

The structure health monitoring sensor responses could be due to the environmental or operational conditions and were not necessarily related to damage or faults. Two types of time series, material temperature and tilt, were considered. The datasets were recorded using sensors provided and installed by the ITMSOIL instrumental company in December 2008 according to the best industrial practice. The bridge was monitored by NPL researchers over a period of four years, and data were recorded every 5 mins, except for a few days of special tests. The installation comprised 10 temperature and 8 tilt sensors and provided regular time series from 2009 to 2012. The sensor locations are shown in Figure 4.19. The temperature measurements are obtained using temperature sensors within the Vibrating Wire Arc Weldable Strain Gauge that is used for temperature compensation. The vibrating wire strain sensors consist of a coil assembly, a vibrating wire element and two weldable anchors. The tilt sensors are Electrolevel Sensors for monitoring movement of structures. The sensors were calibrated at installation, and the monitoring system performed well during the project.

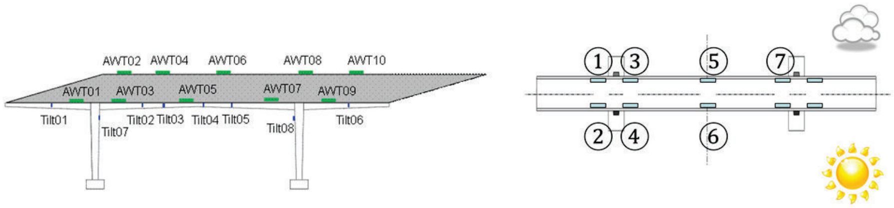

Figure 4.19: Left: material temperature and tilt sensors. Right: orientation of the bridge and temperature sensors towards the Sun.

The experimental program consisted of three types of tests:

1. The 2010- 2011 tests were designed to imitate the consequences of corrosion. Initial stages of concrete spalling were achieved by the removal of concrete around shear links and reinforcing bars (rebars) and the repair of the concrete using materials and techniques according to Highways Agency (HA) specifications. Following this, the rebars on the top of the cantilever were cut to represent the decrease of cross section.
2. The 2011-2012 fatigue and creep tests were carried out on the repaired cantilever and finally the largest loads (about 9 tonnes) were applied twice to approach “component failure” of the repair patches installed on the cantilever. This was the second most significant event of damage introduced.
3. Static and dynamic loading tests were carried out between different stages to assess the performance of the deteriorating structure during the whole period between 2009 and 2012.

The first stage involved removal of the concrete followed by repairs using the best industrial practices and concrete repair materials. Locations and extent of removal reasonably simulated spalling of concrete due to reinforcement corrosion and the deterioration observed in real structures. The repair/strengthening of the cantilever carried out after a period of damage accumulation during various types of tests and was expected to be one of the most significant structural changes. Various static loads were used starting from one tonne applied to one square meter (the serviceability limit for this type of structure) to maximum loads of over nine tonnes over 1.5m^2 during the final stages of the testing program. Fatigue tests used a unique dynamic testing facility, specially designed for this demonstrator by adapting an Instron hydraulic system. About half a million load controlled, 10kN/s cycles were performed.

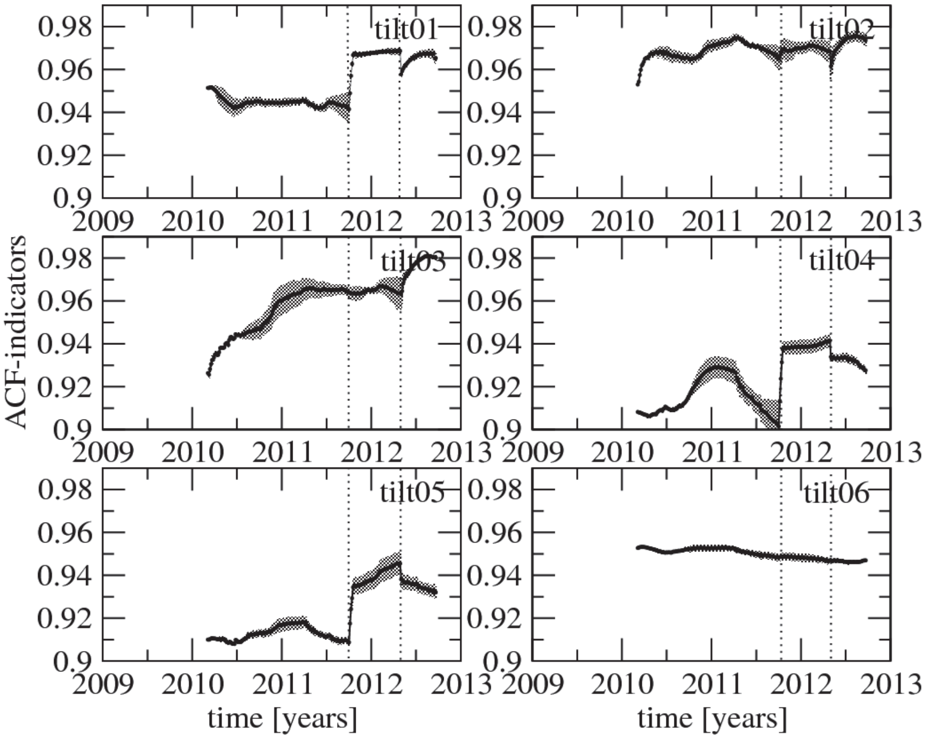

Figure 4.20: Early warning indicators for eight records of tilt sensors. Dashed lines denote the same interval in all six datasets when two abrupt changes occurred.

The recorded sensor datasets have high (5 minutes) and sometimes variable temporal resolution. Hourly equidistant sampling for the purposes of tipping point analysis was chosen. Since the dataset runs for almost four complete years (2009-2012), the datasets have about about 32,000 points. To maintain the same observation interval and temporal resolution, six material temperature and tilt datasets were considered. Temperature was recorded in degrees Celsius and tilt in mm/m, which is equivalent to series recorded in angle radians, because the values are small and tan(angle) approximates angle well. Raw data, deseasonalised data (with subtracted mean annual cycle), and detrended data (with subtracted mean annual cycle and divided by the seasonal cycle of standard deviation) were considered. The raw data had clear seasonal periodicity in both types of sensor data, which should be eliminated, whereas the detrended data appeared as having heavy artefacts of detrending.

The sources of observed artefacts and their influence on complex statistical analysis require further investigation. In this case study, the deseasonalised datasets are analysed. The influence of the bridge experiments that introduced some outliers in the time series was, in terms of the used methods, negligible, because their time scale was up to a few days. The analysis was performed with the datasets where data were thresholded to exclude the oscillations due to the bridge experiments. The results (not shown here) are effectively equivalent to those of the non-thresholded data, because of very small temporal scale of the perturbations.

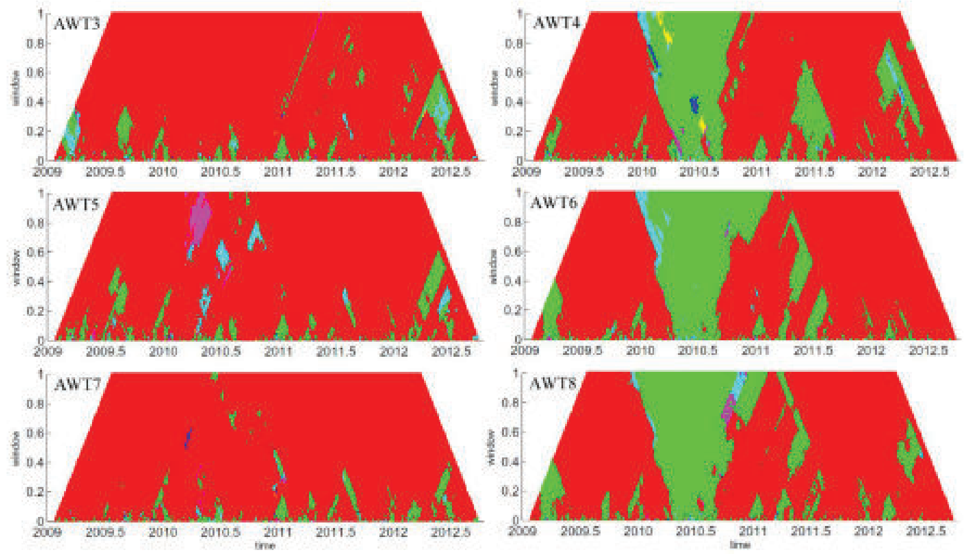

Figure 4.21: Potential analysis plots for six records of material temperature sensors. There is a distinct difference between even (sunny side) and odd (shadow side) parts of the bridge.

Figure 4.21 illustrates the environmental effects on the material conditions of the bridge, as three sensors on the sunny side, during the summer of 2010 [*July's heat wave around the world — World news — guardian.co.uk. London: Guardian. 2010-07-07*, n.d.], demonstrate different patterns compared with the sensors on the shadowy side. This is one of the examples of the environmental signatures of material fluctuations in structure health monitoring.

The exposure of critical transport infrastructure to natural hazards and climate change effects has severe consequences on world economies and societies and, thus, safety and resiliency of transport networks are of paramount importance. The currently available frameworks for quantitative risk and resilience-based design and assessment have been mainly developed for bridges exposed to earthquakes. However, there is an absence of well-informed exposure, vulnerability, functionality and recovery models, which are the main components in the quantification of resilience. Tipping point analysis can serve as a novel tool for the data-driven resilience assessment of transport infrastructure exposed to multiple hazards by using multiscale monitoring data. Monitoring and early warnings are expected to produce accurate and rapidly informed quantitative risk and resilience assessments for transport infrastructure and to enhance asset management.

4.3.2 Electric resistance data

We apply tipping point analysis to measurements of electronic components commonly used in applications in the automotive or aviation industries and demonstrate early warning signals based on scaling properties of resistance time series. The analysis is based on a statistical physics framework with a stochastic model representing the system time series as a composition of deterministic and stochastic components estimated from measurements. The early warning signals are observed much earlier than those estimated from conventional techniques, such as threshold-based failure detection, or bulk estimates used in Weibull failure analysis.

A PCB test vehicle was designed to enable the evaluation of test methods to measure the remaining useful life of solder joints. To accelerate the ageing process, the test boards were placed in a thermal cycling chamber cycling from -55^{o}C to 125^{o}C at a rate of 10^{o}C min^{-1} and with 5-minute dwells at the temperature extremes. We use 5-minute dwells during the data collection from thermal cycling to ensure the whole assembly reaches the set temperature. Five minutes is a standard value to use.

The experimental setup monitored three measurement channels (each channel captured data from a 4-point probe measurement setup). Limitations on experimental equipment availability meant that we were required to run three separate experimental runs to collect the full dataset.

When measuring very low values of resistance (e.g. that of a solder joint), using the 4-point probe method is preferable (to remove any contribution to the measured resistance from the measurement leads). This means we can be confident that measured changes in the resistance are due to the degradation of the solder joint.

Thermal cycling induces failures at interfaces due to a mismatch in the coefficient of thermal expansion (CTE) at those interfaces. Therefore, whilst failure of the resistor would affect the outcomes significantly, it is highly unlikely to occur before solder joint failure. Further to this, if a resistor failed during the test, it would be picked up after the test when the resistors are checked to confirm they are still low resistance.

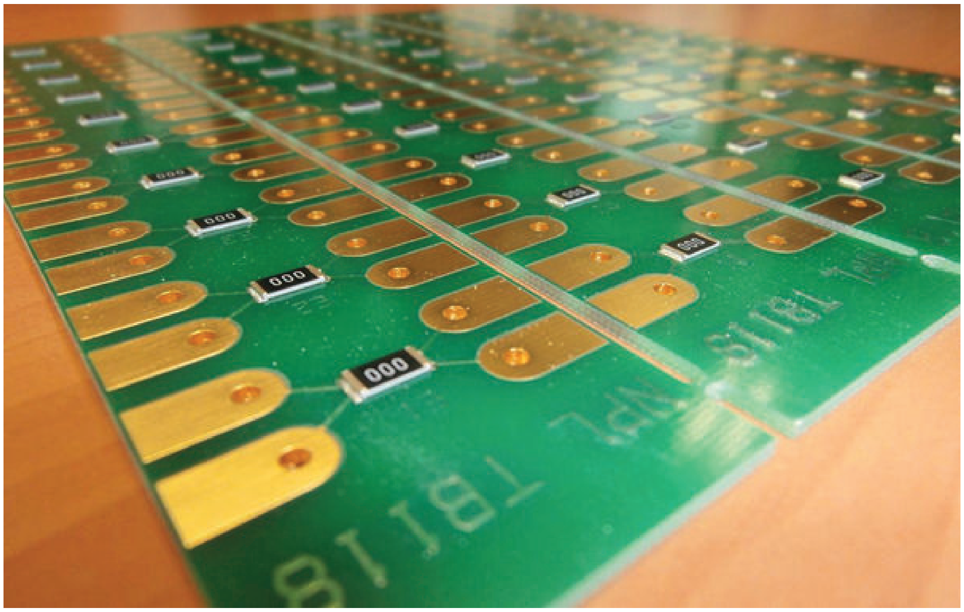

Figure 4.22: Image showing test boards. Each test board could host up to 10 surface mount chip resistors. Each component was connected to four pads to enable a 4-point probe method of resistance monitoring.

The test boards were 1.6 mm thick copper clad FR-4 with a NiAu finish. The test components were zero-ohm 2512 chip resistors connected with a Pb-free solder interconnect. An image of the test board is given in Figure 4.22.

We analysed measured resistance datasets from nine units, which experience failure (critical rise of resistance) after repeated testing cycles, see Figure 4.23. The reported cycles when the units went open circuit: r.1c.1 (run 1 channel 1) — 540, r.1c.2 — 1000, r.1c.3 — 750, r.2c.1 — 1000, r.2c.2 — 815, r.2c.3 — 810, r.3c.1 — 910, r.3c.2 — 543, r.3c.3 — 516.

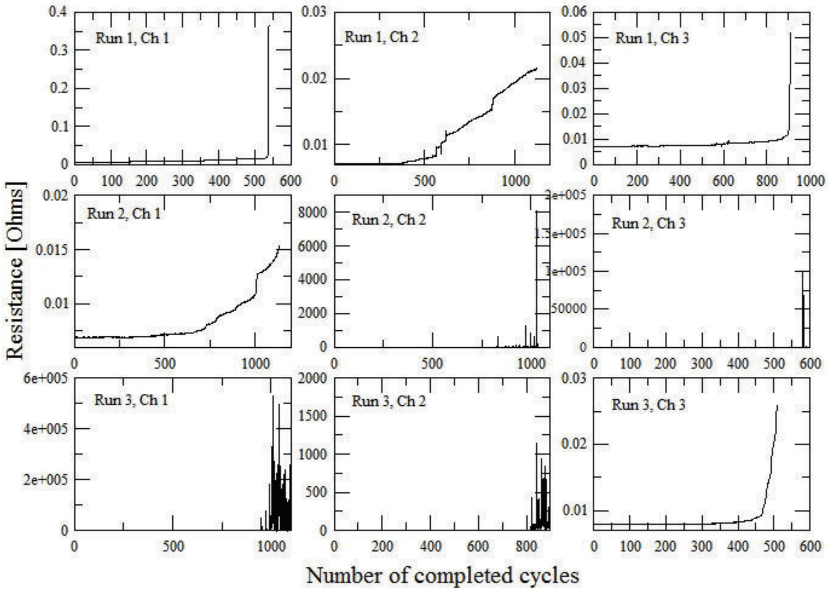

Figure 4.23: Nine time series of measured resistance with critical rise indicating the failure of the units after repeated test cycles (three runs using three channels denoted by 'Ch').

We have calculated two early warning indicators, ACF- and DFA-based, and compared the timing of the obtained early warning signals with the conventional threshold-based warning. As a threshold of failure, one can consider triple-nominal resistance. In the beginning of the experiment, the resistance values of the tested units were about 0.008 Ohm, and therefore the threshold would conventionally be established at about 0.025 Ohm.

We first calculate the ACF-based indicator with uncertainty quantification based on varying window sizes (between 1/4 and 3/4 of the data length) and estimate the time of the early warning signals for them when the ACF-indicator reaches a high value of 0.9, as shown in Figure 4.24. In addition, we consider the average curve of the ACF-indicator and along this curve calculate linear extrapolation of the indicator to estimate when in future it would reach critical value 1 (for DFA, the critical value is 1.5). By doing this, we obtain a set of possible times when the failure would happen, which forms a histogram — this histogram is then used to generate the kernel density of the future failure times. The peak of such a kernel density is the most likely time of failure, statistically. We illustrate this in Figures 4.24,4.25 and also use this information in Figure 4.26.

We assume the real-time situation while moving with sliding windows along the time series and forming the indicators curves (this is what happens when a time series is analysed real-time rather than in retrospect). To estimate the kernel distribution, we perform projections (linear extrapolations) of the indicators curves to obtain the statistics of the future state.

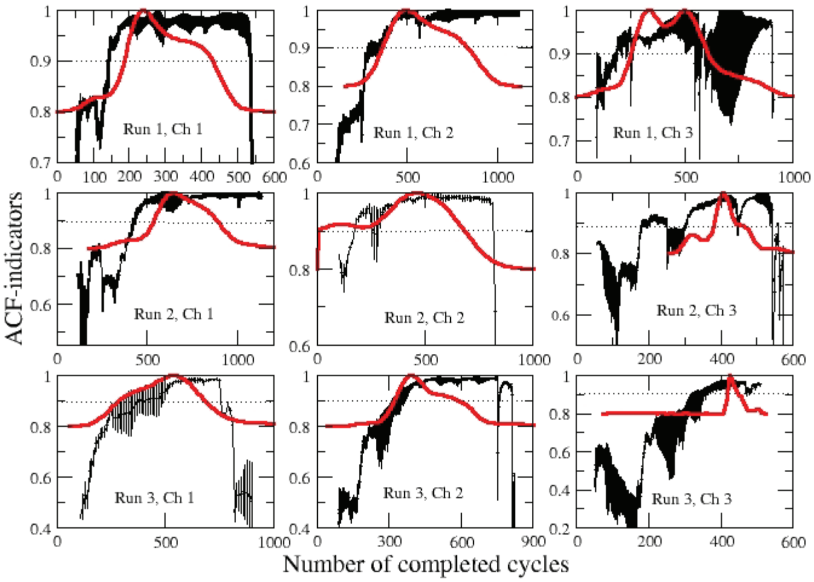

Figure 4.24: ACF-indicators with uncertainty quantification based on varying window size (10-50% of time series length) for resistance data, with clear critical rise prior to failures. Dashed lines denote the high level of auto-correlation. Red curves correspond to probability densities (kernel distributions of the times of projections of subsets of the indicators when critical value 1 would be reached). Panels correspond to the data in Figure 4.23.

We also apply the DFA-indicator to assess early warning signals. In most cases, ACF shows earlier warnings than DFA. This is caused by the difference between single-point ACF estimation (lag-1) and the multiple-point DFA estimate (subset of the DFA curve in the time scale 10-100, as introduced in [Livina and Lenton, 2007]).

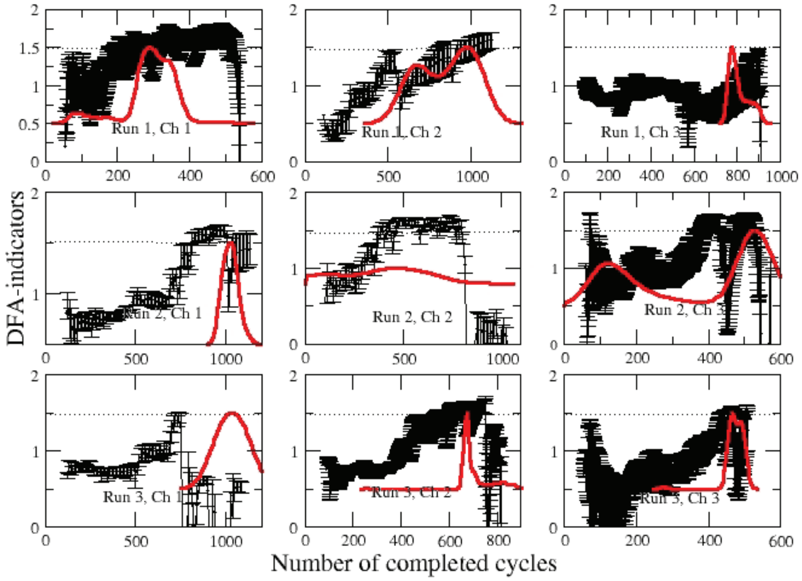

Figure 4.25: DFA-indicators with uncertainty quantification based on varying window size (10-50% of time series length) for resistance data. Dashed lines denote the critical value of the DFA-indicator [Livina and Lenton, 2007]. Red curves correspond to probability densities (kernel distributions of the times of projections of subsets of the indicators when critical value would be reached). Panels correspond to the data in Figure 4.23.

We then map the time points that can be seen in the rising indicators to the plot with the data, in which we also highlight where the electrical interconnect fails and goes open circuit, and observe that early warning signal indicators provide much earlier forewarning than the conventional technique (Figure 4.26). Both early warning signal indicators provide earlier forewarning of the upcoming failure of units with critically rising resistance, as compared with both stringent (magenta arrow) and moderate (red arrow) threshold tests. The stringent test uses the criterion of 20% increase of resistance, whereas the moderate test uses triple-nominal resistance, which is obtained as the mean value of the initial nine estimates of resistance over first 50 cycles. The locations of green and blue arrows are based on the peaks of kernel distributions in Figures 4.24,4.25.

The variability of locations of early warning points in Figure 4.26 is caused by different dynamics of the resistance time series: some of them fail more gradually, whereas others fail abruptly. Most likely, this is related to the material composition of the devices, which vary at mesoscopic level.

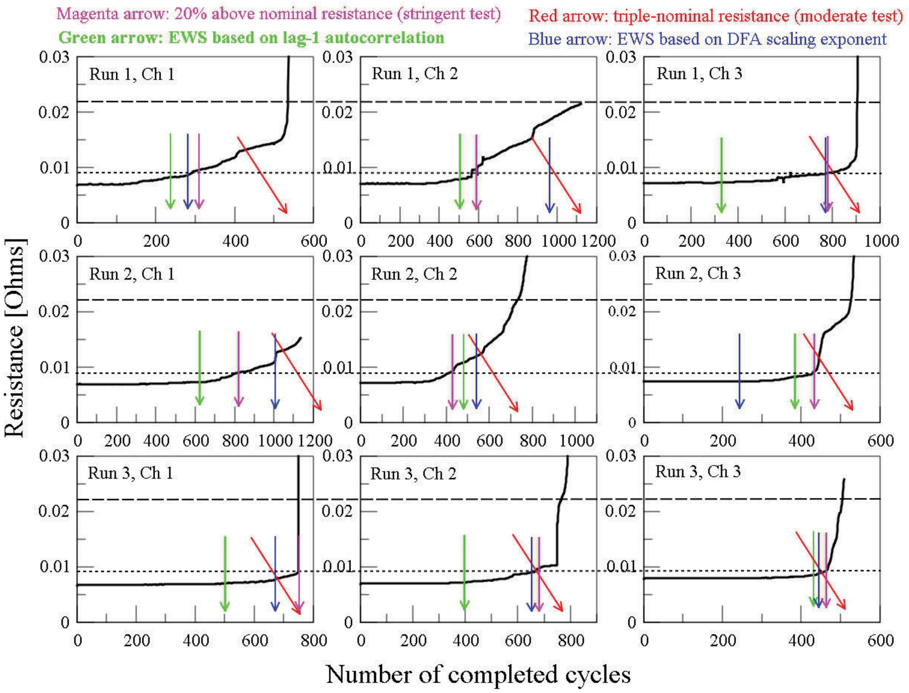

Figure 4.26: Comparison of performance of two early warning signal indicators (green arrow for time stamp — lag-1 ACF-indicator; blue arrow — DFA-indicator) and two conventional test thresholds (magenta arrow for time stamp — stringent test with 20% resistance increase, denoted by dotted line; red arrow — moderate test with triple-nominal resistance threshold, denoted by dashed line). Panels correspond to the data in Figure 4.23.

4.3.3 NASA data of thermal aging of transistors

We consider the NASA experiment of thermal stress published in [Sonnenfeld et al., 2008], where the authors reported on thermal cycling aging experiments conducted on gate-controlled power transistors.

Thermal overstress aging experiments with square signal gate voltage bias were performed in April 2008 with four devices. Device configuration consisted of the emitter connected to the ground of the power supply and the collector with a resistor in series to the positive lead of the power supply. The gate was driven by a high-speed amplifier of the output of a function generator. The experiment monitored the collector current, collector voltage, gate voltage and package temperature of the device. By altering temperature within a range outside its rated standard, steady state recordings of these variables were taken to test durability.

Our analysis aimed to identify, if possible, early warning signals prior to failure of the experimental device: whether the forewarning signal could be obtained from the monitored variables, or from the sets of transients of the experimental data.

For transient measurements, every 22 seconds a transient recording of 100,000 points was captured. These recordings were taken for the collector emitter current signal, the collector emitter voltage, the gate emitter voltage and gate signal voltage each of each of the four devices.

The experiments were performed on each device for duration of one hour, with monitoring of several variables continuously: supply voltage, collector emitter current, heat sink temperature, package temperature, ambient temperature. In addition, high-resolution transients (100,000 points) of gate-emitter voltage were recorded during the experiment, about every minute of the experiment.

Interestingly, the technique of potential analysis [Livina et al., 2010], too, detected various changes in the supply voltage shown in panel (a) of Figure 4.27 — in Figure 4.28 one can see the changes of the number of potential states, where red colour denotes one-well-potential part of the time series. For more details on this technique, see the methodological paper [Livina et al., 2010], where it was introduced.

The package temperature variable was considered, and multivariate canonical correlation analysis was applied using the following variables: supply voltage, node 1 voltage, node 2 voltage, heat sink temperature, and collector emitter current of Device 2. As shown in Figure 4.29, it is possible to detect the final failure at the end of the experiment by using the available time series with training subsets.

CCA, similarly to the potential analysis, allows one to detect the moment of failure. However, while the continuous time series of the experiments with NASA devices (shown in panels (a) of Figure 4.27) picture the step-changes in voltage and temperature, the transients are more useful in identifying the moment when the device becomes vulnerable. The slope of the transient changes, when monitored during the experiment, serves as early warning signal, as shown in panels (d).

It is possible to use the increase in oscillations in the slope indicator as an early warning signal, whereas the results of hierarchical clustering with DTW measure separate the late transients very close to the failure point, as shown by the numbers of the final transients in panels (c).

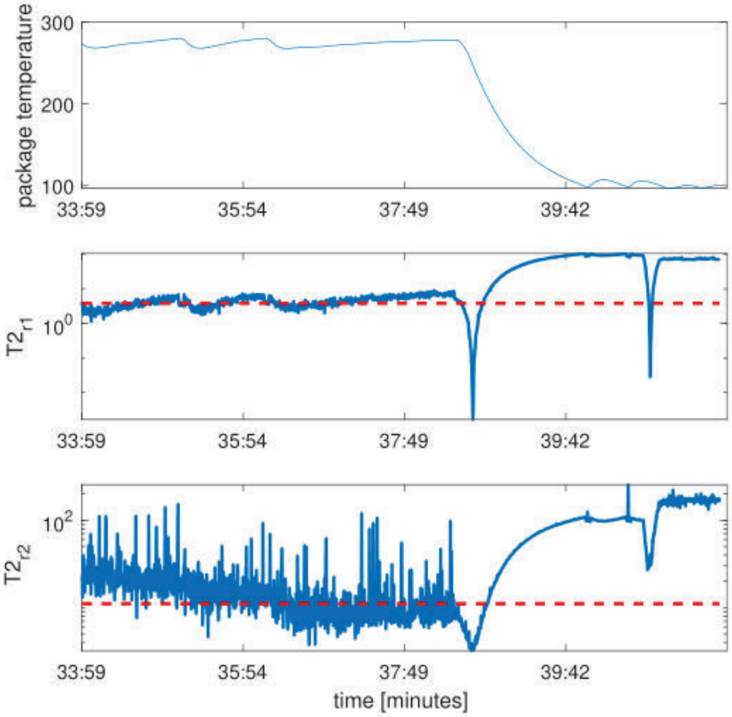

Figure 4.29: Canonical correlation analysis of package temperature variable shown in panel (a) of Figure 4.27 of the NASA thermal aging experiment, Set 8, Device 2. Top panel: fragment of time series with failure; middle and lower panel: CCA statistics with threshold (dashed red line) denoting the failure detection level.

4.3.4 Building management data

Our variable of interest is the supply air temperature (SAT) of a particular zone of the building in the manufacturing plant. SAT is controlled by the average room temperature measured by four sensors installed in the zone using the control signal of pre-stablished fixed setpoint. Average room temperature is then used by the system to regulate the percentage of openness of heating or cooling valves, according to the setpoint. Another factor that influences the temperature of the room is the Outside Air Temperature (OAT). These four variables are represented in Figure 4.30, where the origin, 0, is taken to represent the start of the failure.

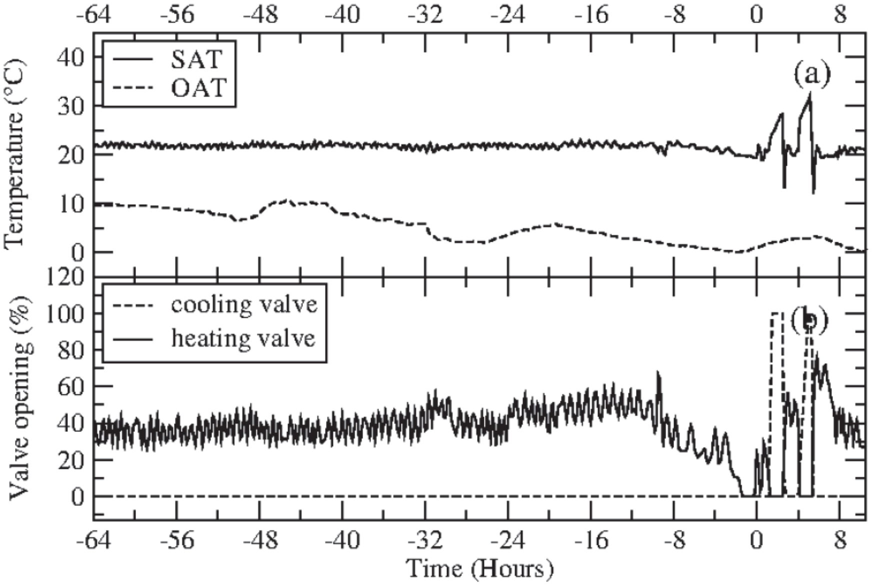

Figure 4.30: Plot of sensor variables affecting room temperature. (a) Supply Air Temperature (SAT) and Outside Air Temperature (OAT) in degrees Celsius. (b) Percentage of opening for each of the heating and cooling valves. Plot window starts at 5:30AM of the 25th of January, 64 hours prior to the failure.

Another variable that controls temperature of the room is the frost coil valve, but we are not using it for the purpose of early failure detection.

- Outside Air Temperature controls frost coil valve. Frost coil valve protects the Air Handling Unit (AHU) from very low temperatures (the parts containing liquid elements, as they can contract and damage the AHU as a consequence).

- Supply Air Temperature controls heating and cooling valves. As data is registered in winter, the cooling valve will be closed most of the time, but in periods of anomaly it activates and provides with cooling to the environment.

We first conduct univariate analysis on the SAT by using different indicators, and then we take into account other variables, reducing their dimensionality and applying the same techniques to compare both approaches and identify which one provides an EWS the earliest.

Results for EWS analysis are presented in Figure 4.31, where SAT is presented together with the EWS indicators. The plot shows data from 64 hours prior to the failure at moment 0, which is when the temperature increases anomalously for the first time. We use different windows to obtain every EWS: 22 hours for variance and ACF1, and 8 hours for PS. The window size has been selected according to the clarity of the signal they provided on each indicator. Y-axis of the variance has been represented in logarithmic scale, for convenience.

The analysis shows that, in general, the indicators provide an earlier and more reliable signal in the multivariate case. This improvement can be seen specially in PS indicator. The difference is that the indicator completely changes its range of movement 24 hours prior to failure. In the multivariate case, the behavior of the system is defined mainly by the heating valve, which steadily opens and closes to provided heating to the environment within an established range or set point.

The early detection of such failures gives time to on-site engineers to make adjustments when necessary before these failures actually happen. This not only reduces maintenance and operational costs, but also produces energy savings by advising when parts of the system should not be activated in some given period, thus compensating the ”blind spots” from the BMS control system by extracting real value from the data generated.

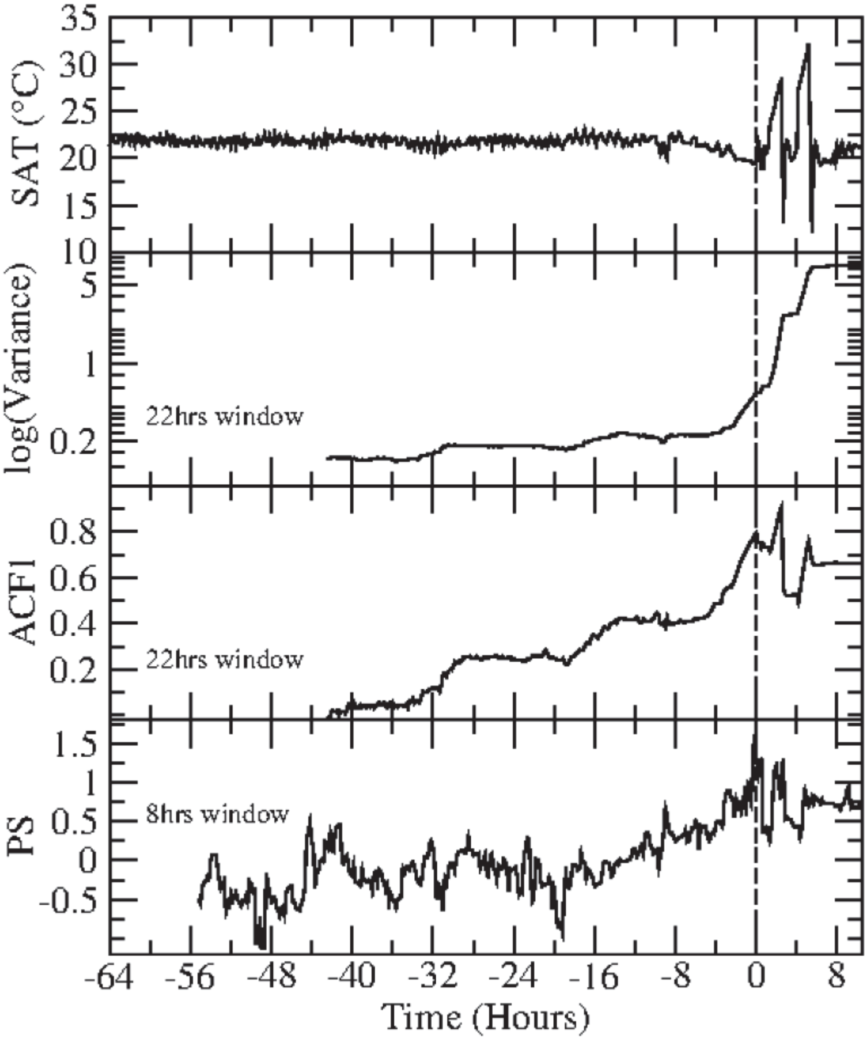

Figure 4.31: EWS analysis of the temperature failure with four indicators: variance, auto-correlation function, and power spectrum. The moment of the failure (when the supply air temperature increases anomalously for the first time) is 9:30AM of the 27th of January 2019. Plot starts at 5:30AM of the 25th of January, 64 hours prior to the failure. Variance is shown in logarithmic scale to make the increase prior to failure more visible. Window sizes used: 22 hours for variance and ACF1, and 16 hours for PS.

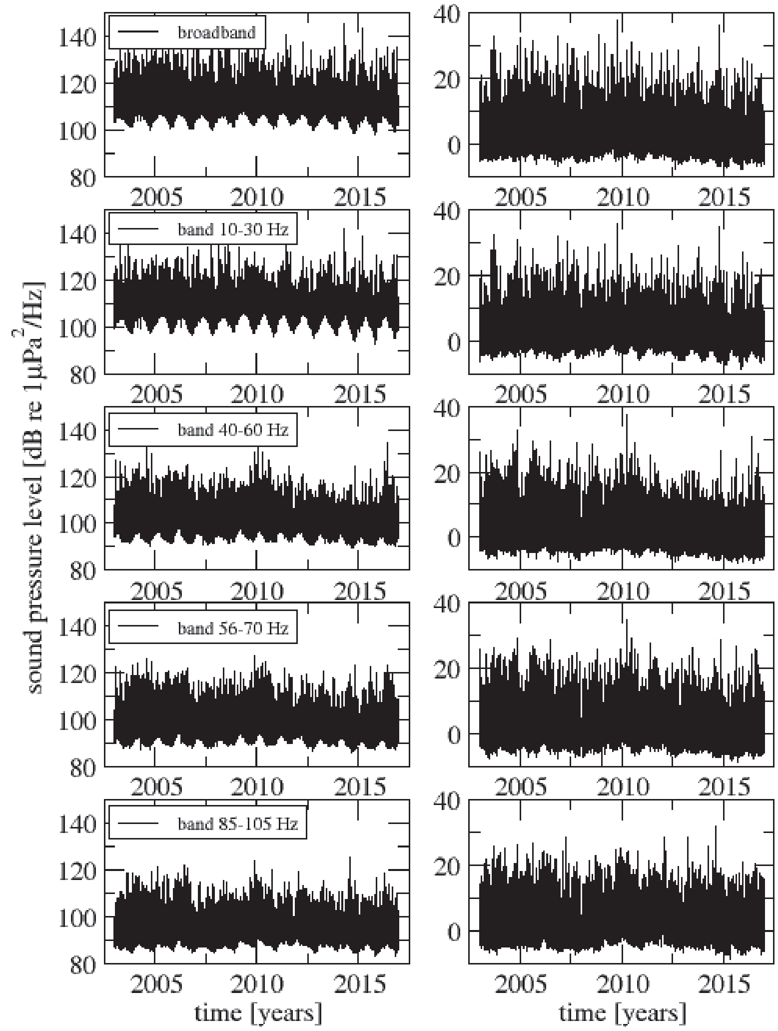

Figure 4.13: Initial (left column) and processed (right column) sound pressure level data in five frequency bands. Processing included interpolation and deseasonalisation. Note that seasonal variability is less pronounced in the higher frequencies of the initial data. At the same time, the records of higher frequencies have a declining trend visible by eye.

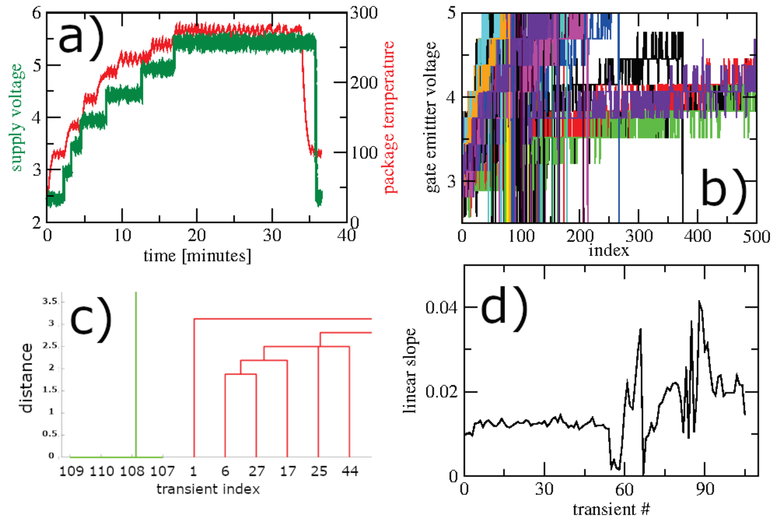

Figure 4.27: NASA thermal aging experiment, Set 8, Device 2: a) combined 2-curve plot of supply voltage and package temperature, indicating the failure at the end of the experiment; b) curves of the of gate emitter voltage transients; c) a fragment of the dendrogram of hierarchical clustering of the transients from panel (b) that shows that the few last transients form a separate cluster with different behaviour; d) slopes of the linear fit of the transients from panel (b) along the experiment that form an early warning indicator.

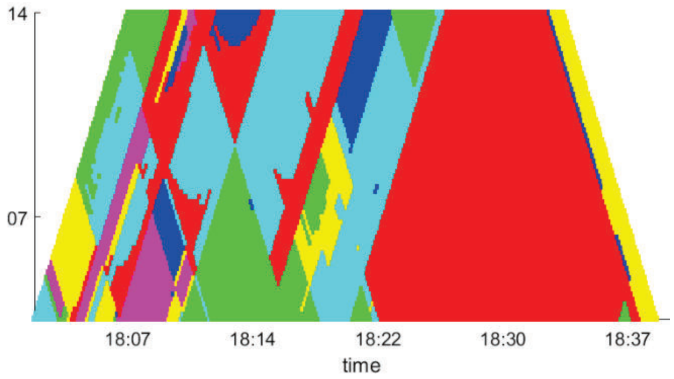

Figure 4.28: Potential plot of the supply voltage variable shown in panel (a) of Figure 4.27 of the NASA thermal aging experiment, Set 8, Device 2. Change of colours denotes different numbers of potential states (i.e., detects bifurcations). These bifurcations signify stable and unstable parts of time series, and the last short period of actual failure, with yellow colour.

Chapter 5

Conclusions

This report studied artificial and observed time series of dynamical systems experiencing transitional and bifurcational tipping, as well as apparent failures.

The critical transitions were detected, forecast, and novel early warning signal indicators were developed for applications in diverse dynamical systems.

In principle, forecasting techniques can be applied at any stage (prior to transition, at the beginning of transition, or when transition has been evolving). However, at earlier stages of transition, because of building up of nonstationarities, the horizon of a successful forecast would be much shorter, which was illustrated in the experiments with artificial data.

It is necessary to mention that forecasting of time series is not always possible or necessary in the case of an abrupt terminal failure: time series may suddenly collapse or critically increase, and forecasting may become obsolete.

As was shown in the recent publication [Billuroglu and Livina, 2022], clustering time series based on patterns may provide novel early warning signals and improve our understanding of dynamical systems. This is an example of a novel technique of early warning signals for anticipation of a fault or tipping, which can be obtained from specific datasets of dynamical systems.

Many of early warning signal indicators are based on temporal scaling in time series that can be quantified using scaling exponents. Interestingly, stochastic models with a Hopf bifurcation are used for developing early warning indicators of tipping, and this brings back the connection with the classical bifurcation theory developed more than a century ago.

There has been a large amount of work involving the use of multiple data sources spread over a geographic area, such as complex network analysis in climate [Gozolchiani et al., 2011, Tsonis et al., 2006, Yamasaki et al., 2008] which has been applied to the El Niño. In [Ludescher et al., 2013, 2014] the cross-covariance is calculated between points on a grid inside a region of interest (in this case, the El Niño basin) and points outside, in contrast with the EOF method which considers the auto-correlation (related to the auto-covariance) of all available points.

The next stage of development of tipping points in climatology was to look at spatial early warning signals within the area of interest [Kefi et al., 2014]. From the initial understanding of a single tipping point in one recorded variable, research efforts moved to studying interdependencies of tipping points and their possible cascades [Dekker et al., 2018], in which tipping points could trigger each other. Still, this was expected to be observed in adjacent tipping elements with obvious dependencies.

A recent study [Liu et al., 2023] applied advanced network analysis to establish long-distance connections (teleconnections) across the globe. The tipping elements considered are Amazon Rainforest Area, the Tibetan Plateau, and West Antarctic ice sheet, and the research has identified strong correlations across long distances between these regions. This is the first time that the theory of complex networks has been applied in the context of tipping points, and the synergy of the two research areas provides an important insight into the global climate dynamics. This work opens a new area of tipping point analysis at global scale. In particular, Liu et al have established the propagation pathway between the Rainforest Amazon Area and West Antarctic Icesheet, and this teleconnection is across almost ten thousand kilometres. The authors explain the path by the steady and strong ocean currents and westerly winds near the West Antarctic area, with transportation of dust and carbonaceous aerosols causing the teleconnection. The network methodology for tipping elements will help detect various teleconnections, which will then be analysed by climatologists to uncover linkage mechanisms, thus providing an exemplary synergy of disciplines.

Recent developments in space-based Earth observations provide rich data for such studies, and it has become possible to analyse large surface regions using reanalysis data. Modelled data (Coupled Model Intercomparison Projects) and reanalysis data, which combines satellite and in-situ observations into global estimates using advanced modelling and data assimilation systems make it possible to systematically study global climatic processes. By combining tipping point analysis and abundant real data, it is possible to conduct global studies in the context of climate change. [Liu et al., 2023] promises high-impact applications in geophysics.

Other directions of research include studying geophysical cascades of tipping points [Dekker et al., 2018], network analysis [Liu et al., 2023], and various applications of Bayesian techniques. The variety of developed tools open broad horizons for development of tipping point analysis for multiple research groups around the globe. The topic of tipping points and bifurcation analysis in applications is experiencing revival with help of numerical methods and recent advances of ML/AI techniques.

The Intergovernmental Panel on Climate Change (IPCC) recognises the growing importance of understanding tipping points and is planning a Special Report on ”Climate Tipping Points and their Implications for Habitability and Resources”, which will be prepared in the framework of the IPCC’s 7th Assessment Cycle, scheduled to start in 2023.

What started some years ago as sparse attempts of analysing early warning signals, nowadays becomes a major research topic in geophysics. This report on tipping point analysis also demonstrates the pathways of applying the statistical physics methodology to a broad range of engineering problems.

Acknowledgements

The report was developed at National Physical Laboratory, funded by the Department for Science, Innovation and Technology (UK). The author is grateful to students and collaborators (in alphabetic order): Dr Stephanie Bell, Berkem Billuroglu, Prof. Alistair Forbes, Dr Peter Harris, Dr Jan Kantelhardt, Dr. Frank Kwasniok, Dr Eng-Tseng Lau, Prof. Tim Lenton, Prof. Gerrit Lohmann, Dr Joaquin Mesa, Dr Joshua Prettyman, Dr Kate Willett.

Bibliography

Abbot, D., Silber, M. and Pierrehumbert, R. [2011], 'Bifurcations leading to summer Arctic sea ice loss', *Journal of Geophysical Research* **116**, D19120.

Amstrup, S. e. a. [2010], 'Greenhouse gas mitigation can reduce sea-ice loss and increase polar bear persistence', *Nature* **468**(7326), 955–958.

Arnold, V. [1986], *Catastrophe Theory*, Springer-Verlag Berlin Heidelberg.

Ashwin, P., Wieczorek, S., Vitolo, R. and Cox, P. [2012], 'Tipping points in open systems: bifurcation, noise-induced and rate-dependent examples in the climate system', *Philosophical Transactions of the Royal Society of London A: Mathematical, Physical and Engineering Sciences* **370**(1962), 1166–1184.

Bak, P., Tang, C. and Wiesenfeld, K. [1988], 'Self-organized criticality', *Physical review A* **38**(1), 364.

Bengtsson, L., Semenov, V. and Johannessen, O. [2004], 'The early twentieth-century warming in the Arctic — a possible mechanism', *Journal of Climate* **17**(20), 4045–4057.

Billuroglu, B. and Livina, V. [2022], 'Full-cycle failure analysis using conventional time series and machine learning techniques', *Journal of Failure Analysis and Prevention* **22**, 1121–1134.

Boe, J., Hall, A. and Qu, X. [2009], 'September sea-ice cover in the Arctic ocean projected to vanish by 2100', *Nature Geoscience* **2**(5), 341–343.

Carlstein, E. [1986], 'The use of subseries values for estimating the variance of a general statistic from a stationary sequence', *The Annals of Statistics* **14**(3), 1171–1179.

Cavalieri, D., Parkinson, C., Gloerson, P. and Zwally, H. [1996], ‘Sea ice concentrations from Nimbus-7 SMMR and DMSP SSM/I passive microwave data, 1978-2007’, `http://nsidc.org/data/nsidc-0051.html`. Accessed on 15.03.2023.

Cimatoribus, A., Drijfhout, S., Livina, V. and van der Schrier, G. [2013], ‘Dansgaard-Oeschger events: bifurcation points in the climate system’, *Climate of the Past* **9**, 323–333.

Comiso, J. [2012], ‘Large decadal decline of the Arctic multiyear ice cover’, *Journal of Climate* **25**, 1176–1193.

Dakos, V., Scheffer, M., van Nes, E. H., Brovkin, V., Petoukhov, V. and Held, H. [2008], ‘Slowing down as an early warning signal for abrupt climate change’, *Proceedings of the National Academy of Sciences* **105**(38), 14308–14312.

Dansgaard, W., Johnsen, J., Clausen, H., Dahl-Jensen, D., Gundestrup, N., Hammer, C., Hvidberg, C., Steffensen, J., Sveinbjörnsdottir, A., Jouzel, J. and Bond, G. [1993], ‘Evidence for general instability of past climate from a 250-kyr ice-core record’, *Nature* **364**, 218–220.

Dekker, M., von der Heydt, A. and Dijkstra, H. [2018], ‘Cascading transitions in the climate system’, *Earth System Dynamics* **9**(12431260).

Ditlevsen, P. D. and Johnsen, S. J. [2010], ‘Tipping points: early warning and wishful thinking’, *Geophysical Research Letters* **37**(19).

Eisenman, I. [2010], ‘Geographic muting of changes in the Arctic sea ice cover’, *Geophysical Research Letters* **37**(16), L16501.

Eisenman, I. and Wettlaufer, J. [2009], ‘Nonlinear threshold behavior during the loss of Arctic sea ice’, *Proceedings of the National Academy of Sciences USA* **106**(1), 28–32.

Feng, M., Meyers, G., Pearce, A. and Wijffels, S. [2003], ‘Annual and interannual variations of the leeuwin current at 32s’, *Journal of Geophysical Research* **108**(11), 3355.

Gardiner, C. [2004], *Handbook of Stochastic Methods*, Springer.

Gilks, W. and Wild, P. [1992], ‘Adaptive rejection sampling for gibbs sampling’, *Applied Statistics - Journal of the Royal Statistical Society Series C* **41**(2), 337–348.

Gladwell, M. [2000], *The Tipping Point: How Little Things Can Make a Big Difference*, Little Brown.

Gozolchiani, A., Havlin, S. and Yamasaki, K. [2011], 'Emergence of El Niño as an autonomous component in the climate network', *Physical review letters* **107**(14), 148501.

Hasselmann, K. [1976], 'Stochastic climate models: Part i. theory', *Tellus* **28**, 473–485.

Held, H. and Kleinen, T. [2004], 'Detection of climate system bifurcations by degenerate fingerprinting', *Geophysical Research Letters* **31**(23).

Heneghan, C. and McDarby, G. [2000], 'Establishing the relation between detrended fluctuation analysis and power spectral density analysis for stochastic processes', *Physical review E* **62**(5), 6103.

Huang, B. e. a. [2014], 'Extended reconstructed sea surface temperature version 4 (ersst.v4)', *Journal of Climate* **28**, 911–930.

ISO18405:2017 [2017], 'Iso underwater acoustics — terminology'.

Jolliffe, I. T. [1986], Principal component analysis and factor analysis, *in* 'Principal component analysis', Springer, pp. 115–128.

Julier, S., Uhlmann, J. and Durrant-Whyte, H. [2000], 'A new method for the nonlinear transformation of means and covariances in filters and estimators', *IEEE Transactions on Automatic Control* **45**, 477–482.

July's heat wave around the world — World news — guardian.co.uk. London: Guardian. 2010-07-07 [n.d.], `https://www.theguardian.com/world/gallery/2010/jul/07/heat-wave`. Accessed on 15.03.2023.

Kantelhardt, J. W., Koscielny-Bunde, E., Rego, H. H., Havlin, S. and Bunde, A. [2001], 'Detecting long-range correlations with detrended fluctuation analysis', *Physica A: Statistical Mechanics and its Applications* **295**(3), 441–454.

Kefi, S., Guttal, V., Brock, W., Carpenter, S., Ellison, A., Livina, V., Seekell, D., Scheffer, M., van Nes, E. and Dakos, V. [2014], 'Early warning signals of ecological transitions: methods for spatial patterns', *PLoS ONE* **9**(3), e92097.

Kleinen, T., Held, H. and Petschel-Held, G. [2003], ‘The potential role of spectral properties in detecting thresholds in the earth system: application to the thermohaline circulation’, *Ocean Dynamics* **53**(2), 53–63.

Koscielny-Bunde, E., Bunde, A., Havlin, H., Eduardo Roman, H., Goldreich, Y. and Schellnhuber, H. [1998], ‘Indication of a universal persistence law governing atmospheric variability’, *Phys. Rev. Lett.* **81**, 729.

Kwasniok, F. [2018], ‘Detecting, anticipating, and predicting critical transitions in spatially extended systems’, *Chaos: An Interdisciplinary Journal of Nonlinear Science* **28**(3), 033614.

Kwasniok, F. and Lohmann, G. [2009], ‘Deriving dynamical models from paleoclimatic records: application to glacial millennial-scale climate variability’, *Physical Review E* **80**(6), 066104.

Kwasniok, F. and Lohmann, G. [2012], ‘A stochastic nonlinear oscillator model for glacial millennial-scale climate transitions derived from ice-core data’, *Nonlinear Processes in Geophysics* **19**(6), 595–603.

Lennartz, S. and Bunde, A. [2009], ‘Eliminating finite-size effects and detecting the amount of whit noise in short records with long-term memory’, *Phys. Rev. E* **79**(066101).

Lenton, T., Livina, V., Dakos, V., Van Nes, E. and Scheffer, M. [2012], ‘Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness’, *Phil. Trans. R. Soc. A* **370**(1962), 1185–1204.

Lenton, T. M., Held, H., Kriegler, E., Hall, J. W., Lucht, W., Rahmstorf, S. and Schellnhuber, H. J. [2008], ‘Tipping elements in the earth’s climate system’, *Proceedings of the national Academy of Sciences* **105**(6), 1786–1793.

Lindsay, R. and Zhang, J. [2005], ‘The thinning of Arctic sea ice, 1988-2003: Have we passed a tipping point?’, *Journal of Climate* **18**, 4879–4894.

Liu, T., Chen, D., Yang, L., Meng, J., Wang, Z., Ludescher, J., Fan, J., Yang, S., Chen, D., Kurths, J., Chen, X., Havlin, S. and Schellnhuber, H. J. [2023], ‘Teleconnections among tipping elements in the earth system’, *Nat. Clim. Chang.* **13**, 67–74.

Livina, V., Barton, E. and Forbes, A. [2014], ‘Tipping point analysis of the npl footbridge’, *Journal of Civil Structural Health Monitoring* **4**, 91–98.

Livina, V., Ditlevsen, P. and Lenton, T. [2012], ‘An independent test of methods of detecting system states and bifurcations in time-series data’, *Physica A* **391**(3), 485–496.

Livina, V., Kwasniok, F. and Lenton, T. [2010], ‘Potential analysis reveals changing number of climate states during the last 60 kyr’, *Climate of the Past* **6**, 77–82.

Livina, V., Kwasniok, F., Lohmann, G., Kantelhardt, J. and Lenton, T. [2011], ‘Changing climate states and stability: from pliocene to present’, *Climate Dynamics* **37**(11-12), 2437–2453.

Livina, V. and Lenton, T. [2007], ‘A modified method for detecting incipient bifurcations in a dynamical system’, *Geophys. Res. Lett.* **34**, L03712.

Livina, V. and Lenton, T. [2013], ‘A recent tipping point in the Arctic sea-ice cover: abrupt and persistent increase in the seasonal cycle since 2007’, *The Cryosphere* **7**, 275–286.

Livina, V., Lewis, A. and Wickham, M. [2020], ‘Tipping point analysis of electrical resistance data with early warning signals of failure for predictive maintenance’, *Journal of Electronic Testing* **36**, 569–576.

Livina, V., Lohmann, G., Mudelsee, M. and Lenton, T. [2013], ‘Forecasting the underlying potential governing the time series of a dynamical system’, *Physica A* **392**(18), 3891–3902.

Ludescher, J., Gozolchiani, A., Bogachev, M. I., Bunde, A., Havlin, S. and Schellnhuber, H. J. [2013], ‘Improved El Niño forecasting by cooperativity detection’, *Proceedings of the National Academy of Sciences* **110**(29), 11742–11745.

Ludescher, J., Gozolchiani, A., Bogachev, M. I., Bunde, A., Havlin, S. and Schellnhuber, H. J. [2014], ‘Very early warning of next El Niño’, *Proceedings of the National Academy of Sciences* **111**(6), 2064–2066.

Makse, H. A., Havlin, S., Schwartz, M. and Stanley, H. E. [1996], ‘Method for generating long-range correlations for large systems’, *Physical Review E* **53**(5), 5445.

Maslanik, J. and Stroeve, J. [1999], ‘Near-real-time dmsp ssmis daily polar gridded sea ice concentrations, version 1’, `https://nsidc.org/data/nsidc-0081/versions/1`. Accessed on 15.03.2023.

Meier, W., Stewart, J., Wilcox, H., Hardman, M. and Scott, D. [2021], 'Near-real-time dmsp ssmis daily polar gridded sea ice concentrations, version 2', `https://nsidc.org/data/nsidc-0081/versions/2`. Accessed on 15.03.2023.

Milnor, J. [1963], *Morse Theory*, Princeton University Press.

Mudelsee, M. [2012], *Climate Time Series Analysis: Classical Statistical and Bootstrap Methods*, Springer.

Neagoe, V. [1990], 'Chebyshev nonuniform sampling cascaded with the discrete cosine transform for optimum interpolation', *IEEE Transactions on Acoustics, Speech and Signal Processing* **38**(10), 1812–1815.

Nghiem, S. e. a. [2007], 'Rapid reduction of Arctic perennial sea ice', *Geophysical Research Letters* **34**, L19504.

Notz, D. [2009], 'The future of ice sheets and sea ice: Between reversible retreat and unstoppable loss', *Proceedings of the National Academy of Sciences USA* **106**, 20590–20595.

Palmer, T. and Williams, P. [2008], 'Stochastic physics and climate modelling', *Phil. Trans. R. Soc. A.* **366**, 2419–2425.

Perry, M., Livina, V. and Niewczas, P. [2016], 'Tipping point analysis of cracking in reinforced concrete', *Smart Materials and Structures* **25**(1), 015027.

Petoukhov, V. and Semenov, V. [2010], 'A link between reduced barents-kara sea ice and cold winter extremes over northern continents', *J. Geophys. Res.* **115**(D21), D21111.

Poincare, H. [1892], *Les Methodes Nouvelles de la Mecanique Celeste*, Vol. 1, Gauthier-Villars, Paris.

Prettyman, J., Kuna, T. and Livina, V. [2018], 'A novel scaling indicator of early warning signals helps anticipate tropical cyclones', *Europhysics Letters* **121**, 10002.

Robinson, S., Lepper, P. and Hazelwood, R. [2014], 'Good practice guide for underwater noise measurement. npl good practice guide no. 133'.

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. [2009], 'Early-warning signals for critical transitions', *Nature* **461**(7260), 53–59.

Scheffer, M., Carpenter, S., Foley, J. A., Folke, C. and Walker, B. [2001], ‘Catastrophic shifts in ecosystems’, *Nature* **413**(6856), 591–596.

Semenov, V., Park, W. and Latif, M. [2009], ‘Barents sea inflow shutdown: A new mechanism for rapid climate changes’, *Geophysical Research Letters* **36**(14), L14709.

Sherman, M., Speed, F. and Speed, F. [1998], ‘Analysis of tidal data via the blockwise bootstrap’, *Journal of Applied Statistics* **25**(3), 333–340.

Silverman, B. [1986], *Density estimation of statistics and data analysis*, Chapman & Hall.

Sitz, A., Schwarz, U., Kurths, J. and Voss, H. [2002], ‘Estimation of parameters and unobserved components for nonlinear systems from noisy time series’, *Physical review E* **66**(1), 016210.

Sonnenfeld, G., Goebel, K. and Celaya, J. [2008], ‘An agile accelerated aging, characterization, and scenario simulation system for gate controlled power transistors’.

Stroeve, J., Holland, M., Meier, W., Scambos, T. and Serreze, M. [2007], ‘Arctic sea ice decline: faster than forecast’, *Geophysical Research Letters* **34**(L09501).

Stroeve, J., Kattsov, V., Barrett, A., Serreze, M., Pavlova, T., Holland, M. and Meier, W. [2012], ‘Trends in Arctic sea ice extent from cmip5, cmip3 and observations’, *Geophys. Res. Lett.* **39**(L16502).

Surovyatkina, E., Kravtsov, Y. and Kurths, J. [2005], ‘Fluctuation growth and saturation in nonlinear oscillators on the threshold of bifurcation of spontaneous symmetry breaking’, *Phys Rev E* **72**(4), 046125.

Svensson, A., Andersen, K., Bigler, M., Clausen, H., Dahl-Jensen, D., Davies, S., Johnsen, S., Muscheler, R., Parrenin, F., Rasmussen, S., Ruthlisberger, R., Seierstad, T., Steffensen, J. and Vinther, B. [2008], ‘A 60 000 year greenland stratigraphic ice core chronology’, *Climate of the Past* **4**, 47–57.

Thompson, J. and Sieber, J. [2011], ‘Predicting climate tipping as a noisy bifurcation: a review’, *International Journal of Bifurcation and Chaos in Applied Sciences and Engineering* **21**(2), 399–423.

Tietsche, S., Notz, D., Jungclaus, J. and Marotzke, J. [2011], 'Recovery mechanisms of Arctic summer sea ice', *Geophysical Research Letters* **38**(2), L02707.

Trenberth, K. and Caron, J. [2000], 'The southern oscillation revisited: Sea level pressures, surface temperatures, and precipitation', *Journal of Climate* **13**, 4358–4365.

Tsonis, A. A., Swanson, K. L. and Roebber, P. J. [2006], 'What do networks have to do with climate?', *Bulletin of the American Meteorological Society* **87**(5), 585–596.

Van Nes, E. H. and Scheffer, M. [2007], 'Slow recovery from perturbations as a generic indicator of a nearby catastrophic shift', *The American Naturalist* **169**(6), 738–747.

Veraart, A. J., Faassen, E. J., Dakos, V., van Nes, E. H., Lürling, M. and Scheffer, M. [2012], 'Recovery rates reflect distance to a tipping point in a living system', *Nature* **481**(7381), 357–359.

von Storch, H. and Zwiers, F. [2002], *Statistical Analysis in Climate Research*, Cambridge University Press, Cambridge.

Welch, P. [1967], 'The use of fast fourier transform for the estimation of power spectra: a method based on time averaging over short, modified periodograms', *IEEE Transactions on audio and electroacoustics* **15**(2), 70–73.

Wiesenfeld, K. [1985], 'Noisy precursors of nonlinear instabilities', *Journal of Statistical Physics* **38**, 1071–1097.

Willett, K., Dunn, R., Thorne, P., Bell, S., Podesta, M., Parker, D., Jones, P. and Williams, C. [2014], 'Hadisdh land surface multi-variable humidity and temperature record for climate monitoring', *Clim. Past* **10**, 1983–2006.

Williamson, M. S. and Lenton, T. M. [2015], 'Detection of bifurcations in noisy coupled systems from multiple time series', *Chaos: An Interdisciplinary Journal of Nonlinear Science* **25**(3), 036407.

Winton, M. [2006], 'Does the Arctic sea ice have a tipping point?', *Geophysical Research Letters* **33**(23), L23504.

Yamasaki, K., Gozolchiani, A. and Havlin, S. [2008], ‘Climate networks around the globe are significantly affected by El Niño’, *Physical review letters* **100**(22), 228501.

Yokoyama, Y., Lambeck, K., De Deckker, P., Johnston, P. and Fifield, L. [2000], ‘Timing of the last glacial maximum from observed sea-level minima’, *Nature* **406**, 713–716.

Yue, S., Pilon, P. and Cavadias, G. [2002], ‘Power of the mann–kendall and spearman’s rho tests for detecting monotonic trends in hydrological series’, *Journal of hydrology* **259**(1), 254–271.

Zhao, K., Wulder, M., Hu, T., Bright, R., Wu, Q., Qin, H., Li, Y., Toman, E., Mallick, B., Zhang, X. and Brown, M. [2019], ‘Detecting change-point, trend, and seasonality in satellite time series data to track abrupt changes and nonlinear dynamics: A bayesian ensemble algorithm’, *Remote Sensing of Environment* **232**, 111181.

Publications on tipping points

1. Livina V., Network analysis: Connected climate tipping elements, *Nature Climate Change* **13**, 15-16 (2023).
2. Billuroglu B. and V. Livina, Full-cycle failure analysis using conventional time series and machine learning techniques, *Journal of Failure Analysis and Prevention* **22**, 1121-1134 (2022)
3. Prettyman J. et al, Power spectrum scaling as a measure of critical slowing down and precursor to tipping points in dynamical systems, *Environmental Research Letters* **17**, 035004 (2022)
4. Mesa J. et al, Early Warning Signals of Failures in Building Management Systems, *International Journal of Metrology and Quality Engineering* **12**, 11 (2021).
5. Argyroudis S. et al, Data-driven resilience assessment for transport infrastructure exposed to multiple hazards by integrating multiscale terrestrial and airborne monitoring systems, *Proceedings of IABMAS: 10th International conference on Bridge Maintenance, Safety and Management*, 2021.
6. Livina V. et al, Tipping point analysis of electrical resistance data with early warning signals of failure for predictive maintenance, *Journal of Electronic Testing* **36**, 569-576 (2020); arXiv:1904.04636
7. Prettyman J. et al, Generalised early warning signals in multivariate and gridded data, *Chaos* **29**, 073105 (2019).
8. Prettyman J. et al, A novel scaling indicator of early warning signals helps anticipate tropical cyclones, *Europhysics Letters* **121**, 10002 (2018)
9. Livina V. et al, Tipping point analysis of ocean acoustic noise, *Nonlinear Processes in Geophysics* **25**, 89-97 (2018).
10. Perry M. et al, Tipping point analysis of cracking in reinforced concrete, *Smart Materials and Structures* **25 (1)**, 015027 (2016).
11. Lenton T. and V. Livina, Detecting and anticipating climate tipping points, in *”Extreme events: observations, modelling, and economics”* (Eds. M.Chavez, M.Ghil, and J.Urrutia-Fucugauchi), AGU monograph 214, Wiley-Blackwell, pp. 51-62 (2016)

12. Livina V. et al, Tipping point analysis of atmospheric oxygen concentration, *Chaos* **25**, 036403 (2015). The pdf-preprint is available in ResearchGate.

13. Kefi S. et al, Early warning signals of ecological transitions: methods for spatial patterns, *PLoS ONE* **9 (3)**, e92097 (2014).

14. Livina V. et al, Tipping point analysis of the NPL footbridge, *Journal of Civil Structural Health Monitoring* **4**, 91-98 (2014).

15. Drijfhout S. et al, Spontaneous, abrupt climate change due to an atmospheric blocking/sea-ice/ocean feedback in an unforced climate model simulation, *Proceedings of the National Academy of Sciences* **110 (49)**, 19713-19718 (2013).

16. Livina V. et al, Forecasting the underlying potential governing the time series of a dynamical system, *Physica A* **392 (18)**, 3891-3902 (2013).

17. Cimatoribus A. et al, Dansgaard-Oeschger events: bifurcation points in the climate system, *Climate of the Past* **9**, 323-333 (2013).

18. Livina V. and T. Lenton, A recent tipping point in the Arctic sea-ice cover: abrupt and persistent increase in the seasonal cycle since 2007, *The Cryosphere* **7**, 275-286 (2013).

19. Dakos V. et al, Methods for Detecting Early Warnings of Critical Transitions in Time Series Illustrated Using Simulated Ecological Data, *PLoS ONE* **7(7)**, e41010 (2012).

20. Lenton T. et al, Climate bifurcation during the last deglaciation? *Climate of the Past* **8**, 1127-1139 (2012).

21. Lenton T. et al, Early warning of climate tipping points from critical slowing down: comparing methods to improve robustness, *Philosophical Transactions of the Royal Society A* **370**, 1185-1204 (2012).

22. Livina V. et al, An independent test of methods of detecting system states and bifurcations in time-series data, *Physica A* **391 (3)**, 485-496 (2012).

23. Livina V. et al, Changing climate states and stability: from Pliocene to present, *Climate Dynamics* **37 (11-12)**, 2437-2453 (2011).

24. Vaz Martins T. et al, Resonance induced by repulsive interactions in a model of globally coupled bistable systems, *Phys. Rev. E* **81**, 041103 (2010).

25. Livina V. et al, Potential analysis reveals changing number of climate states during the last 60 kyr, *Climate of the Past* **6**, 77-82 (2010).

26. Lenton T. et al, Using GENIE to study a tipping point in the climate system, *Philosophical Transactions of the Royal Society A* **367**, 871 (2009).

27. Livina V. and T. Lenton, A modified method for detecting incipient bifurcations in a dynamical system, *Geophys. Res. Lett.* **34**, L03712 (2007).