

ENSO dynamics in current climate models: an investigation using nonlinear dimensionality reduction

I. Ross¹, P. J. Valdes¹, and S. Wiggins²

¹School of Geographical Sciences, University of Bristol, University Road, Bristol BS8 1SS, UK

²School of Mathematics, University of Bristol, University Walk, Bristol BS8 1TW, UK

Received: 9 October 2007 – Revised: 18 February 2008 – Accepted: 26 March 2008 – Published: 22 April 2008

Abstract. Linear dimensionality reduction techniques, notably principal component analysis, are widely used in climate data analysis as a means to aid in the interpretation of datasets of high dimensionality. These linear methods may not be appropriate for the analysis of data arising from nonlinear processes occurring in the climate system. Numerous techniques for nonlinear dimensionality reduction have been developed recently that may provide a potentially useful tool for the identification of low-dimensional manifolds in climate data sets arising from nonlinear dynamics. Here, we apply Isomap, one such technique, to the study of El Niño/Southern Oscillation variability in tropical Pacific sea surface temperatures, comparing observational data with simulations from a number of current coupled atmosphere-ocean general circulation models. We use Isomap to examine El Niño variability in the different datasets and assess the suitability of the Isomap approach for climate data analysis. We conclude that, for the application presented here, analysis using Isomap does not provide additional information beyond that already provided by principal component analysis.

1 Introduction

The El Niño/Southern Oscillation (ENSO) is the most important mode of interannual variability in the Earth's climate, driven by atmosphere-ocean interactions in the equatorial Pacific, but with effects reaching as far as north-eastern North America and Europe (Philander, 1990; McPhaden et al., 2006). ENSO events (El Niño and La Niña) occur on an irregular basis at intervals of 2–7 years, and individual ENSO events are variable in their evolution and effects. As a recent example, the 1997/1998 El Niño exhibited behaviour not previously seen, with westerly wind bursts linked to the Madden-Julian Oscillation playing an important role; some

observers have suggested that this may represent a transition to a new regime of ENSO behaviour (McPhaden, 1999; Fedorov and Philander, 2000).

The representation of ENSO in climate models is of interest because of the long-range effects of ENSO on the climate system, both around the Pacific and further afield. Recent model intercomparison studies considering ENSO variability in the current generation of coupled atmosphere-ocean general circulation models (AchutaRao and Sperber, 2002, 2006; van Oldenborgh et al., 2005; Guilyardi, 2006) indicate that some aspects of ENSO variability are represented better in current models than in earlier generations of GCMs (Neelin et al., 1992; Latif et al., 2001), including the overall frequency of El Niño events and enhanced temperature variability over the eastern Pacific. However, current models still display significant deficiencies in the representation of the ocean-atmosphere coupling mechanisms important for ENSO variability – see particularly van Oldenborgh et al. (2005) on this point, where the individual feedback mechanisms relating wind stress, thermocline depth and sea surface temperature are examined in detail in current models.

There has been considerable disagreement about the exact source of ENSO variability. One point of view is that ENSO arises from unstable modes of variability in the tropical ocean-atmosphere system, with limits to predictability determined by growth in errors in initial conditions associated with chaotic dynamics (e.g., Zebiak and Cane, 1987; Jin et al., 1994; Tziperman et al., 1994). The other possibility is that ENSO is a damped linear oscillation excited by stochastic forcing, the limits to predictability being inherent in the stochastic nature of the forcing (e.g., Burgers, 1999; Moore and Kleeman, 1999; Thompson and Battisti, 2000). One deficiency of linear models is that they are not able to reproduce the observed asymmetry between El Niño and La Niña events. This asymmetry has previously been investigated using measures based on sea surface temperature (SST) variance and skewness, nonlinear dynamical heating and explicit characterisation of symmetric and asymmetric structures in

Correspondence to: I. Ross
(ian@skybluetrades.net)

SSTs during ENSO events (An, 2004; An and Jin, 2004; An et al., 2005a; Monahan and Dai, 2004). These studies reveal wide variations in the representation of ENSO asymmetry in coupled ocean-atmosphere models.

Independent of the exact mechanism of variability, the spatial coherence of ENSO in the Pacific leads us to expect that there should be a low-dimensional model that captures at least some of the variability in the tropical ocean-atmosphere system. Here, we approach the assessment of ENSO in coupled GCMs by attempting to identify such low-dimensional structures in the dynamics of the tropical Pacific atmosphere and ocean. It should be noted that, in general, the mechanisms leading to ENSO and ENSO-like variability in current coupled atmosphere-ocean GCMs show significant differences compared to the mechanisms contributing to ENSO variability in the real atmosphere-ocean system. For instance, van Oldenborgh et al. (2005) report that most of the models that they examine show a response of the zonal wind field to equatorial SST anomalies that is weaker and more confined to equatorial latitudes than seen in observations. This weak wind response is compensated by a stronger direct response of SSTs to changes in the wind field and a weaker damping of SST variations than observed. This different balance of factors in the models compared to the observations should lead us to view conclusions drawn from models about ENSO variability in the real atmosphere and ocean with some caution. However, it is still of interest to examine how well we can characterise what low-dimensional dynamics is seen in the models, and to see if this characterisation can provide any further insight into the behaviour of the models. For instance, earlier studies have indicated that ENSO variability can be approximated as a two-dimensional oscillation, one degree of freedom being associated with the NINO3 SST index, the mean SST anomaly across the region 150° W–90° W, 5° S–5° N, and the other with the equatorial Pacific warm water volume (Burgers, 1999; Kessler, 2002; McPhaden, 2003). These two degrees of freedom vary in approximate quadrature during El Niño events. One would hope that any analysis method aimed at characterising ENSO variability in observational or simulated data would be able to identify these two degrees of freedom.

Our question here is, given high-dimensional data from observations or model simulations, what is the best way to characterise low-dimensional behaviour? We are interested in attempting to infer low-dimensional dynamics from relatively limited amounts of data. Observational time series from the Pacific provide around 100 years of monthly SSTs. Time series of several hundred years are available from coupled GCM simulations. To facilitate inter-model comparison, we wish to proceed in a “black box” fashion, adopting a purely data-driven approach without using information about the internal features of the models we are studying.

The method most commonly used in climate data analysis for this type of dimensionality reduction is principal component analysis (PCA) (von Storch and Zwiers, 2003), which

uses an eigendecomposition of the input data covariance matrix. In the problem considered here, we analyse a time series of n maps of sea surface temperature, each with m ocean points. Discarding non-ocean points in each map, we use SST measurements from the remaining ocean data points to construct data vectors $\mathbf{x}_i \in \mathbb{R}^m$, with $i=1, \dots, n$. The covariance matrix of this SST data is then

$$\mathbf{C} = \langle (\mathbf{x} - \langle \mathbf{x} \rangle)(\mathbf{x} - \langle \mathbf{x} \rangle)^T \rangle,$$

where $\langle \cdot \rangle$ denotes time averaging. We write the eigenvector decomposition of \mathbf{C} as $\mathbf{C} = \mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$, with $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_m)$ a diagonal matrix of the eigenvalues λ_i in descending order of magnitude, and \mathbf{Q} a matrix whose columns are the corresponding eigenvectors \mathbf{q}_i . The eigenvectors $\mathbf{q}_i \in \mathbb{R}^m$ are spatial patterns of variation in the data, often called empirical orthogonal functions (EOFs). The first of these, \mathbf{q}_1 , represents the direction in data space with the greatest variance, \mathbf{q}_2 the direction orthogonal to \mathbf{q}_1 with the next greatest variance in the data, and so on. The time series of SST maps, \mathbf{x}_i , can then be expanded in terms of the orthogonal basis provided by the EOFs as

$$\mathbf{x}_i = \sum_j \alpha_{ij} \mathbf{q}_j.$$

The coefficients α_{ij} are called the principal component (PC) time series and give the temporal variation in the data in each of the orthogonal directions in data space spanned by the EOFs. The eigenvalue associated with each EOF measures the proportion of the total variance of the input data explained by that EOF. With the EOFs in descending eigenvalue order, we may extract an EOF subset explaining some pre-selected proportion of the total variance, $V_p = \{\mathbf{q}_i | 1 \leq i \leq p\}$ say, where p is the number of EOFs required to explain the required proportion of the total variance. By projecting the input data into the subspace $\mathcal{V}_p = \text{span}(V_p)$, we arrive at a reduced dimensionality representation of variability in the input data. Compared to the original data this reduced representation has the minimum squared error totalled over all data points of any choice of projection basis of dimension p .

The primary disadvantage of PCA for our purposes is that it is only able to project into linear subspaces of the original m -dimensional data space. If our data points, instead of lying in a linear subspace, lie in a curved low-dimensional submanifold of the data space, PCA will generally not detect the full structure of the data manifold, instead approximating it by the nearest linear subspace in a least squares sense.

This limitation has led to the development of a wide range of nonlinear dimensionality reduction approaches. Of these schemes, only a small number have previously been applied to ENSO data. These include both the method used in this study (Tenenbaum et al., 2000; Gámez et al., 2004) and methods based on neural networks, either self-organising maps (Leloup et al., 2007) or multilayer perceptrons (Monahan, 2001; An et al., 2005b; Wu and Hsieh, 2003).

The neural network approach that has seen most application to questions of ENSO variability is nonlinear principal component analysis (NLPCA), described by Monahan (2001) and, including extensions to canonical correlation analysis and singular spectral analysis, Hsieh (2004). The NLPCA method uses an autoassociative neural network having input and output layers with numbers of neurons corresponding to the number of dimensions of the input data (a preliminary projection into the space spanned by the first few EOFs is normally used to reduce the dimensionality of the input data without losing a significant amount of the data variance), hidden layers attached to the input and output layers with as many neurons as required to give a good fit to the input data, and a “bottleneck layer” between the two hidden layers, whose architecture determines the form of the reduced dimensionality data produced. The neural network is trained on the data set whose dimensionality is to be reduced, the weights in the network being varied so as to reduce the mean squared error between the input data (applied to the input layer of the network) and the network outputs. The idea is thus to produce a network reproducing the input data as faithfully as possible, with information passing through the bottleneck layer, which has a restricted number of neurons. The outputs of the neurons in the bottleneck layer are then taken to be a reduced dimensionality representation of the input data. A single neuron in the bottleneck layer produces a one-dimensional reduced representation of the input data, two neurons in the bottleneck layer a two-dimensional reduced representation, and so on. Additional constraints can be imposed on the structure of the bottleneck layer to yield reduced representations with required characteristics. The most obvious example of this is a “circular” bottleneck node, with two degrees of freedom whose values are constrained to define a point on the unit circle. This yields a bottleneck layer representing a periodic one-dimensional system. Application of NLPCA with one- and two-dimensional bottleneck layers to tropical Pacific observational SST data demonstrated that low-dimensional NLPCA approximations characterise variability in the data better than the corresponding linear PCA approximations, and that NLPCA approximations are able to represent the asymmetry between El Niño and La Niña seen in the observational data (Monahan, 2001). Application of NLPCA with a circular bottleneck layer to observational thermocline depth data from the equatorial Pacific successfully captured the oscillatory nature of thermocline depth variations through the ENSO cycle, and identified differences in the behaviour of the recharge and discharge phases of the oscillation (An et al., 2005b). Further applications of NLPCA in the context of studies of ENSO include the use of nonlinear canonical correlation analysis (NLCCA) to identify nonlinear correlations between SST and wind stress variations in the equatorial Pacific (Wu and Hsieh, 2003).

Many other nonlinear dimensionality reduction techniques have been developed, mostly in the machine vision community, to address issues of feature identification and mo-

tion tracking (e.g., Bishop et al., 1998; Roweis and Saul, 2000; Broomhead and Kirby, 2005; Hinton and Salakhutdinov, 2006; Lin et al., 2006). Many of these methods can be placed into a common framework along with PCA by considering them as seeking a transformation that preserves “interesting” geometric information in the input data. In the case of PCA, this “interesting” information is the Euclidean distances between data points; the required transformation is thus a simple linear orthogonal transformation. A more complex example is the algorithm used in this study, Isomap (Tenenbaum et al., 2000). Isomap finds a nonlinear transformation that preserves not Euclidean distances between data points, but an approximation to distances between data points as measured along geodesics in the data manifold. These geodesic distances are an intrinsic feature of the dynamics of the system under study and are not dependent on the details of the embedding in the observation space. Further elaborations of this idea are possible. For instance, Lin et al. (2006) introduce a method they call Riemannian Manifold Learning, which attempts to preserve not only an approximation to geodesic distances in the data manifold, but also local curvature information. We do not consider this method further in this study.

The only previous application of Isomap to climate data analysis of which we are aware is the work of Gámez et al. (2004), where Isomap was applied to observational SSTs for the equatorial Pacific to examine ENSO variability. Gámez et al.’s results are substantially replicated by our raw SST analysis of the NOAA ERSST v2 observational dataset (Sect. 5.1) and we extend their analysis to consider results from coupled atmosphere-ocean GCMs. As well as being of intrinsic interest, ENSO variability provides a good test case for nonlinear dimensionality reduction methods, primarily because the expected results are relatively easy to interpret. ENSO is by far the strongest mode of climate variability after the annual cycle and has both a clear signature of temporal variability and easily recognisable spatial patterns.

The plan of the paper is as follows. In Sect. 2, we describe the datasets we use. In Sect. 3, we present some conventional analyses of ENSO behaviour in the model simulations to set the scene for interpretation of the Isomap results. In Sect. 4, we describe the Isomap algorithm and examine some issues relating to the sensitivity of Isomap to tunable parameters in the algorithm. Section 5 presents results from performing Isomap analyses on tropical Pacific SST datasets. Finally, in Sect. 6, we present conclusions and recommendations concerning the use of Isomap in climate data analysis.

2 Data and models

In this study, we examine ENSO variability in tropical Pacific SST data from a variety of observational and model sources. Since ENSO is a coupled ocean-atmosphere phenomenon, it would be better to examine other variables in

conjunction with SST, in particular thermocline depth and surface wind stress. However, our goal here is to perform a simple inter-model comparison, so we initially restrict our analysis to SST. We include some analysis of thermocline depth variations later, mostly in the form of equatorial warm water volume.

As observational SST data, we use the NOAA ERSST v2 dataset (Smith and Reynolds, 2004). This is a global dataset running from 1854 to the present day at $2^\circ \times 2^\circ$ resolution, constructed from SST observations using statistical reconstructions in regions with sparse observations. Because of a lack of observations in the equatorial Pacific before about 1900, most variability in this region in the early part of the time series is due solely to the annual cycle. For the purposes of this study, we extract a 100-year subset from 1900–1999 of the full ERSST v2 time series in order to reduce problems due to non-stationarity. There is still some residual non-stationarity in the SST observations associated with changes in ENSO behaviour over time, but this is small. We also examined another observational SST dataset covering a comparable period, the UK Meteorological Office HadISST 1.1 dataset (UK Meteorological Office, 2006). Results were similar to those reported here. In Sects. 5.1 and 5.2, we examine correlations between Isomap results and equatorial Pacific warm water volume (WWV) time series. For observational WWV data, we use the time series derived by Meinen and McPhaden (2000).

Model simulations from a range of coupled ocean-atmosphere GCMs were used for this study, utilising results from the World Climate Research Programme's (WCRP) Coupled Model Intercomparison Project phase 3 (CMIP3) multi-model dataset (Table 1). In this study, we use data from pre-industrial control simulations (`picntrl`) in the CMIP3 database. We do not use all of the CMIP3 models, excluding from consideration simulations that show little or no interannual tropical Pacific atmosphere-ocean variability, either because of the model structure or due to other unidentified problems (e.g., the GISS-AOM and GISS-ER models). For all model simulations, monthly SST time series are used, the length of the time series available for each model being shown in Table 1. Warm water volume time series were calculated for all models where ocean body temperature data was available by determining the depth of the 20°C isotherm by linear interpolation, then integrating the volume of water above the 20°C isotherm in the region 120°E – 80°W , 5°S – 5°N , as in McPhaden (2003).

3 Model ENSO behaviour

Before considering Isomap, we present some more conventional analyses of ENSO behaviour. First, we consider the climatology and magnitude of interannual variability of equatorial Pacific SSTs. Figure 1a shows annual mean SST in the equatorial Pacific, averaged between 2°S and 2°N . Al-

though most of the models show a cold bias across the Pacific basin, with SSTs up to 4°C cooler than observed, they do simulate the gradient of mean SST from the Western Warm Pool around Indonesia to the cooler waters of the eastern Pacific. However, most of the models do not show a monotonic eastwards decline in SST across the basin, instead exhibiting an upturn in mean SST from 100 – 120°W to the eastern edge of the basin. These warmer temperatures near the eastern basin boundary have been observed in previous inter-model comparisons of tropical Pacific SST variability (Mehcho et al., 1995; Latif et al., 2001; AchutaRao and Sperber, 2002) and have been ascribed to difficulties in modelling marine stratus clouds in this region, the steep orography near the coast of South America and the narrow coastal upwelling zone in the eastern Pacific. It appears that relatively little progress has been made in correcting this deficiency in current coupled GCMs.

Figure 1b shows the annual standard deviation of SST across the Pacific in the same latitude band. Here, observations show low variability in the western Pacific and higher variability in the east, where conditions vacillate between the normal cold tongue state and El Niño conditions, characterised by the incursion of warmer water from the western Pacific into the east. Some of the models represent this pattern reasonably well, although the gradient in variability is represented less well than the gradient in mean SST, and again there are problems for all of the models at the far eastern end of the Pacific basin, probably for the same reasons as for the mean SST. The range of variability of the modelled SSTs is quite wide, with one model (FGOALS-g1.0) showing variability as much as 2.5 times the observed values. Some models (CGCM3.1(T47), CGCM3.1(T63), MIROC3.2(hires) and MIROC3.2(medres)) simulate essentially no gradient in variability across the basin.

The SST variability data displayed in Fig. 1b can be summarised using the NINO3 SST index. High values of this index reflect El Niño conditions and low values La Niña conditions. The fifth column of Table 1 shows the standard deviation of the NINO3 SST index for each of the models used here. For comparison, the standard deviation of NINO3 SST for the ERSST v2 observational dataset is 1.26°C for the period 1900–2000. The results in Table 1 indicate that most of the models have a reasonable range of NINO3 SST variability, with CGCM3.1(T47), CGCM3.1(T63) and UKMO-HadGEM1 having too little variability and CNRM-CM3 and FGOALS-g1.0 too much. (As noted above, a number of other models in the CMIP3 model inter-comparison were not used in this study because of unrealistically low NINO3 SST variability. Only models with a NINO3 SST standard deviation of 0.5°C or greater were included in this study.)

The temporal variability of ENSO can be examined using power spectra of the NINO3 SST anomaly time series. Figure 2 shows such spectra calculated using a maximum entropy method (Press et al., 1992, Sect 13.7). The observations show a broad and low peak for periods between about

Table 1. Models used in this study, atmosphere and equatorial ocean spatial resolutions, lengths of simulation available (L), NINO3 SST index standard deviations (σ_{NINO3}), number of ocean grid points in the region 125°W – 65°W , 20°S – 20°N (m), line style used in later plots and references to model documentation. Model horizontal resolution is expressed as degrees longitude \times degrees latitude or a spectral grid designation and vertical resolution as L_n , where n is the number of model levels.

| Model | Atmosphere resolution | Ocean resolution | L (yr) | σ_{NINO3} ($^\circ\text{C}$) | m | Legend | Reference |
|------------------|-------------------------------------|-------------------------------------|----------|--|-------|---------|------------------------------|
| BCCR-BCM2.0 | T63 L31 | $1.5^\circ \times 0.5^\circ$ L35 | 250 | 1.44 | 6133 | — | (Furevik et al., 2003) |
| CCSM3 | T85 L26 | $1.125^\circ \times 0.27^\circ$ L40 | 500 | 1.06 | 19550 | — | (Collins et al., 2006) |
| CGCM3.1(T47) | T47 L31 | $1.85^\circ \times 1.85^\circ$ L29 | 500 | 0.59 | 1742 | — | (Kim et al., 2002) |
| CGCM3.1(T63) | T63 L31 | $1.4^\circ \times 0.94^\circ$ L29 | 400 | 0.64 | 4473 | — | (Kim et al., 2002) |
| CNRM-CM3 | T63 L45 | $2^\circ \times 0.5^\circ$ L31 | 430 | 1.90 | 3049 | — | (Salas-Méla et al., 2005) |
| CSIRO-Mk3.0 | T63 L18 | $1.875^\circ \times 0.84^\circ$ L31 | 380 | 1.26 | 3395 | — | (Gordon et al., 2002) |
| ECHO-G | T30 L19 | $2.75^\circ \times 0.5^\circ$ L20 | 341 | 1.51 | 3418 | - - | (Min et al., 2005) |
| FGOALS-g1.0 | T42 L26 | $1^\circ \times 1^\circ$ L33 | 350 | 1.98 | 6281 | - - | (Yu et al., 2004) |
| GFDL-CM2.0 | $2.5^\circ \times 2^\circ$ L24 | $1^\circ \times 1/3^\circ$ L50 | 500 | 1.37 | 10073 | - - | (Delworth et al., 2006) |
| GFDL-CM2.1 | $2.5^\circ \times 2^\circ$ L24 | $1^\circ \times 1/3^\circ$ L50 | 500 | 1.52 | 10073 | - - | (Delworth et al., 2006) |
| GISS-EH | $5^\circ \times 4^\circ$ L20 | $2^\circ \times 2^\circ$ L16 | 400 | 1.03 | 6172 | - - | (Schmidt et al., 2006) |
| INM-CM3.0 | $5^\circ \times 4^\circ$ L21 | $2^\circ \times 2.5^\circ$ L33 | 330 | 1.29 | 1276 | - - | (Volodin and Diansky, 2004) |
| IPSL-CM4 | $2.5^\circ \times 3.75^\circ$ L19 | $2^\circ \times 1^\circ$ L31 | 500 | 1.19 | 3078 | - - - - | (Marti et al., 2005) |
| MIROC3.2(hires) | T106 L56 | $0.28^\circ \times 0.187^\circ$ L47 | 100 | 1.20 | 9944 | - - - - | (K-1 model developers, 2004) |
| MIROC3.2(medres) | T42 L20 | $1.4^\circ \times 0.5^\circ$ L43 | 500 | 1.14 | 6527 | - - - - | (K-1 model developers, 2004) |
| MRI-CGCM2.3.2 | T42 L30 | $2.5^\circ \times 0.5^\circ$ L23 | 350 | 1.06 | 2583 | - - - - | (Yukimoto et al., 2006) |
| UKMO-HadCM3 | $3.75^\circ \times 2.5^\circ$ L19 | 1.25° L20 | 341 | 1.13 | 3926 | - - - - | (Gordon et al., 2000) |
| UKMO-HadGEM1 | $1.875^\circ \times 1.25^\circ$ L38 | $1^\circ \times 1/3^\circ$ L40 | 240 | 0.97 | 11337 | - - - - | (Johns et al., 2006) |

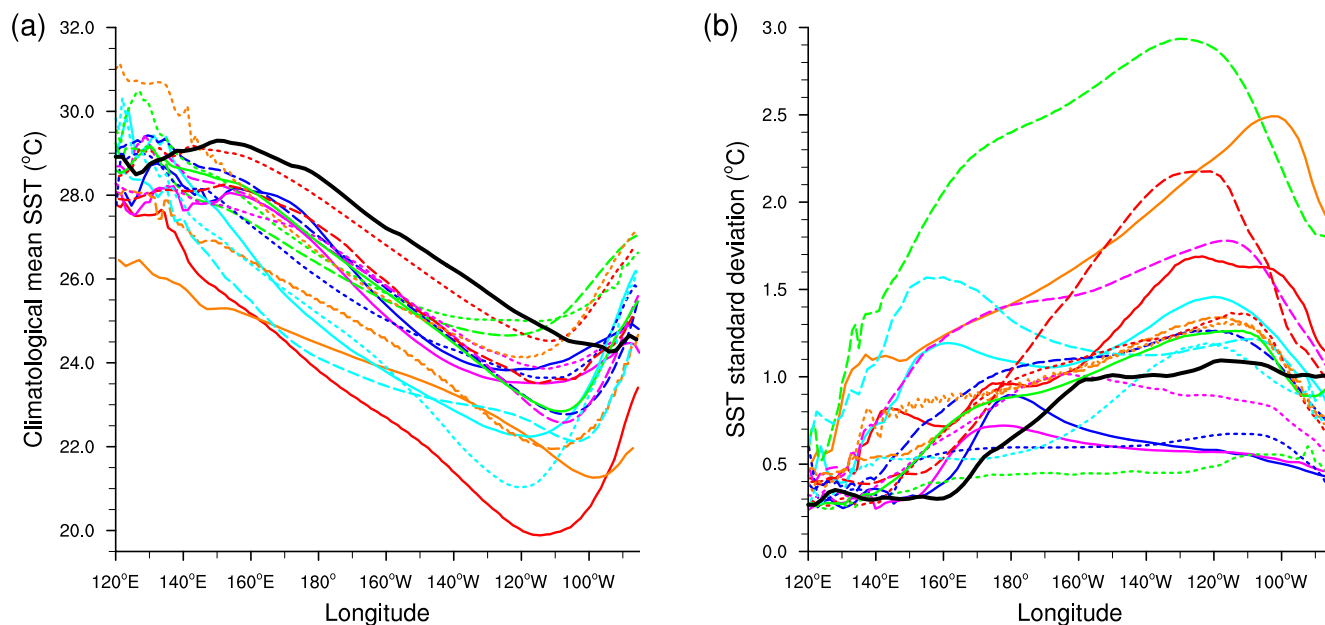


Fig. 1. Climatological mean SST (a) and annual standard deviation of SST (b) across the equatorial Pacific from observations (thick black line) and models (coloured lines – see Table 1 for key). Values shown are averaged between 2°S and 2°N .

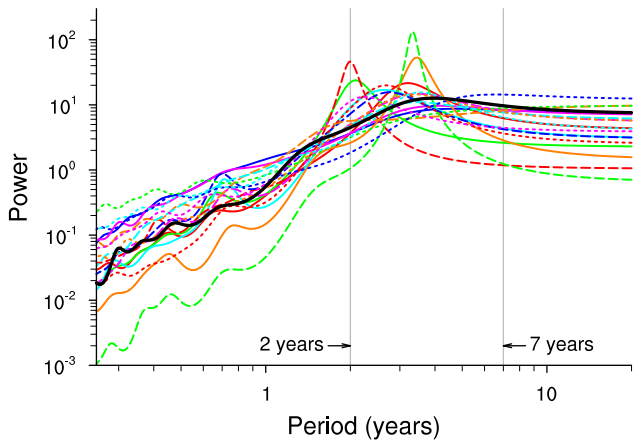


Fig. 2. Maximum entropy power spectra of NINO3 SST index variability from observations (thick black line) and models (coloured lines – see Table 1 for key). All spectra are calculated using 20 poles.

2 and 7 years, indicating the temporal irregularity of ENSO. Among the models, this pattern is replicated most closely in the GFDL-CM2.1, INM-CM3.0 and UKMO-HadCM3 simulations. Other models show either weaker variability in the ENSO band, or variability that is too strongly peaked around a single frequency. This is particularly evident for CCSM3, CNRM-CM3, ECHO-G and FGOALS-g1.0. For the more extreme of these models, one can question whether these narrowband signals can really be identified with ENSO, since they lack the characteristic broad power spectrum of observed ENSO variability.

A common linear dimensionality reduction approach used for geophysical and climate data analysis is principal component analysis (PCA) (von Storch and Zwiers, 2003), also known as empirical orthogonal function (EOF) analysis. The relationship between this linear method and the nonlinear Isomap method will be explored in Sect. 4, but here we present some PCA results for our SST datasets. We calculated area weighted EOFs and principal component time series for SST anomalies from all datasets across the region 125°W – 65°W , 20°S – 20°N . For the observed ERSST v2 SSTs, we used data for the period 1900–2000, while for the models, we used all of the available output, with simulation lengths as listed in Table 1. In each case, after computation, the EOFs are normalised to have unit maximum amplitude for ease of plotting. The corresponding principal component time series are rescaled accordingly.

The first three EOFs from the observations are shown in Figs. 3a–c. The first EOF (Fig. 3a) shows a SST pattern substantially similar to that of a fully developed El Niño event, with warmer temperatures stretching across the equatorial Pacific, replacing the normal tongue of cooler water lying in the eastern Pacific. This first EOF explains 53.2% of the total

variance in the SST data. The second EOF (Fig. 3b) shows a northwest-southeast oriented dipole pattern centred around 120°W , 10°N , explaining about 9.5% of the total observed variance, while the third EOF (Fig. 3c) explains 8.3% of the data variance and shows an east-west dipole lying along the equator with centres of action at around 160°W and near the coast of South America.

These patterns of observed spatial variability can be compared to results from the model simulations. Some selected results are shown in Figs. 3d–o. Here, we display the first three EOFs for CCSM3 (Figs. 3d–f), FGOALS-g1.0 (Figs. 3g–i), GFDL-CM2.1 (Figs. 3j–l) and UKMO-HadCM3 (Figs. 3m–o). The patterns seen represent a cross-section of the typical behaviour seen in the models. In each case, the first EOF is of approximately the right shape, but stretches too far to the west across the Pacific. In the observed data, the region of greatest weight in the first EOF lies well to the east of the date line, while in the model results it extends westwards to 150°E or further. Also, only the pattern for GFDL-CM2.1 has a reasonable shape in the far eastern sector of the Pacific, with the other models either having a pattern not properly connected to South America (CCSM3 and UKMO-HadCM3), or with too much spread of the EOF pattern near the western coast of South and Central America (FGOALS-g1.0). In addition, the range of total variance explained by the first EOF differs quite widely between the models. CCSM3 (explained variance of 46.7%) and GFDL-CM2.1 (explained variance of 60.2%) are closest to the range seen in the observational data, while FGOALS-g1.0 (77.7%) and UKMO-HadCM3 (39.3%) lie outside the observed range, reflecting the unrealistically high (FGOALS-g1.0) and low (UKMO-HadCM3) ENSO variability seen in the NINO3 SST index in these models (Table 1, column 5). The second and third EOFs from the model simulations present a less clear picture. Their spatial patterns are variable; CCSM3 and UKMO-HadCM3 both display a second EOF bearing some resemblance to that of the observational data, with a northwest-southeast dipole centred at about 145°W , 5°N , but the second EOF pattern in FGOALS-g1.0 is more complex, and that seen in GFDL-CM2.1 has a distinct equatorial dipole pattern, more similar to the third EOF of the observational data than to the second. There is great variability in the pattern of the second EOF seen in the other models (not shown).

In principal component analysis, the EOFs represent the spatial patterns of different modes of variability (for real EOFs, actually standing oscillations), while temporal variability is captured in the principal component (PC) time series. Each PC time series gives the projection of the input data time series onto its corresponding EOF, and because of the orthogonality of the EOFs, the PC time series are all linearly uncorrelated by construction. Despite this lack of linear correlation, there are clear nonlinear relationships between the PC time series in the Pacific SST datasets examined here. This can be seen in Fig. 4, which shows selected scatter plots

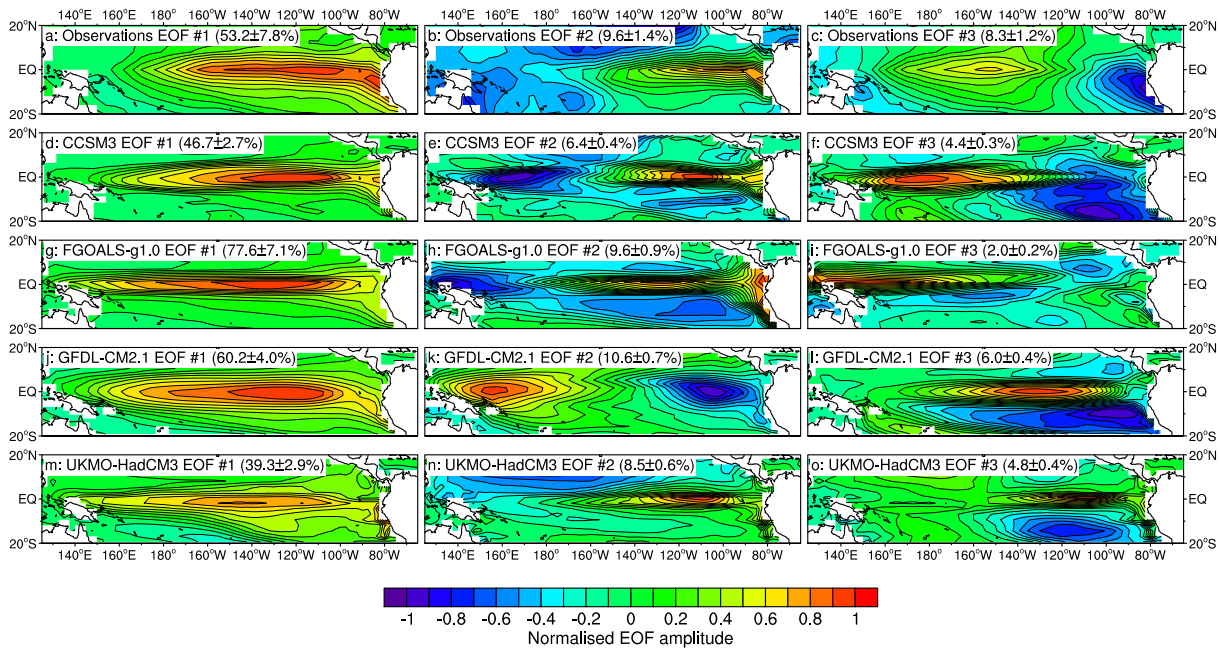


Fig. 3. Sea surface temperature EOFs for the ERSST v2 observational dataset (a–c), CCSM3 (d–f), FGOALS-g1.0 (g–i), GFDL-CM2.1 (j–l) and UKMO-HadCM3 (m–o). Each EOF is normalised to have unit maximum amplitude. Explained variance for each EOF is shown in parentheses, with 95% confidence intervals calculated using the asymptotic results quoted in Hannachi et al. (2007).

of PC time series values. Figure 4a shows PC#1 plotted versus PC#2 for the observational ERSST v2 dataset. Although the two PC time series are not linearly correlated, the asymmetry in the PC scatter plot indicates that they may not be truly independent, and that there may be a nonlinear relationship between the values of PC#1 and PC#2, with large positive and negative values of PC#1 being associated with larger positive values of PC#2. This is because, on average, warm anomalies along the equator east of 150° W during El Niño events are of greater magnitude than cold anomalies during La Niña events. This relationship has previously been discussed in the context of applying other nonlinear dimensionality reduction methods to Pacific SST data (Monahan, 2001). Similar, and in some cases, even stronger, nonlinear relationships are seen between the PC time series for model SSTs. Figure 4b shows a scatter plot of PC#1 and PC#2 from the UKMO-HadCM3 model. Here, there is a similar asymmetric pattern to that seen in the observational data. Although it is difficult to ascribe this to any specific physical mechanism in the model, it seems likely that the root of the asymmetry is similar to that seen in the observations. Whatever the origin of the relationship, the scatter plot is not the Gaussian cloud that would be expected if the PC time series were derived from a simple linear process. Similar comments can be made about the more extreme nonlinearity displayed in Fig. 4c, a scatter plot of PC#1 versus PC#2 for GFDL-CM2.1. This is particularly striking because GFDL-CM2.1 is one of the models from the CMIP3 ensemble that has the

most realistic ENSO (van Oldenborgh et al., 2005). Here, the greater asymmetry in the PC scatter plot may be partially due to the wide meridional spread of the spatial pattern of the first SST EOF and the very distinct zonal dipole pattern in the second SST EOF. Similarly nonlinear PC#1 versus PC#2 scatter plots are seen for some other models with similar structures in their first two EOFs (GFDL-CM2.0 and ECHO-G and, to a lesser extent, MRI-CGCM2.3.2). Any mechanistic explanation of this nonlinearity would require a more detailed analysis of the different ocean-atmosphere feedbacks in the GFDL-CM2.1 model, along the lines of (van Oldenborgh et al., 2005).

The analyses presented so far could be characterised as “conventional” approaches to climate data analysis. Through these analyses, we see a wide range of behaviour in the models, corresponding more or less closely to the behaviour seen in observations. It appears that PCA may not be the most appropriate tool to use here, primarily because of the strong nonlinear relationships between the different PC time series derived from the data.

4 The Isomap algorithm

4.1 Algorithm description

The Isomap algorithm is a two-step process that simultaneously attempts to find a low-dimensional manifold on which a set of data points lies, and Euclidean coordinates giving




Fig. 4. Scatter plots of SST PC #1 versus PC #2 for ERSST v2 observations (a), UKMO-HadCM3 (b) and GFDL-CM2.1 (c).

the locations of the data points in this low-dimensional manifold. The first step in the algorithm is to use a graph-based approximation to the data manifold to calculate approximate geodesic distances between the data points (Sect. 4.1.1). These geodesic distances are then analysed using multidimensional scaling (MDS) to find a Euclidean embedding of the data manifold (Sect. 4.1.2).

4.1.1 Geodesic approximation

As will be explained below, PCA can be considered as an application of the same multidimensional scaling approach used in Isomap, but employing a Euclidean distance function. Isomap uses a distance function that approximates geodesic distances in the data manifold. The aim of this is to determine the intrinsic structure of the data manifold without the more rigid constraints that come from using Euclidean distances.

Geodesics in the data manifold are approximated in two stages. First, a weighted graph is constructed whose vertices are the data points and whose edges connect each point to its nearest neighbours, as determined by Euclidean distances between the data points. The edge weights are the Euclidean distances. There are two ways of setting up this nearest neighbour graph. A distance threshold, ε , can be used, so that edges are included in the graph from a point to all other points closer than ε . If the set of points is denoted by V , the nearest neighbour graph \mathcal{G}_ε is then

$$\mathcal{G}_\varepsilon = (V, E_\varepsilon) = (V, \{(x, y) \mid x, y \in V, d_E(x, y) < \varepsilon\}),$$

where $d_E(x, y)$ is the Euclidean distance between points x and y . The main benefit of this definition is that it is somewhat insensitive to inhomogeneities in data point sampling, and can lead to more robust MDS results. Its primary disadvantage is that it is difficult to establish a reasonable value for ε without some experimentation and it may be necessary to select an inappropriately large value for ε in order to ensure that the graph \mathcal{G}_ε is connected. The second approach is to use a nearest neighbour count, k , so that the nearest neighbour graph contains, for each data point, edges to the k

nearest neighbours. The graph \mathcal{G}_k is then defined as

$$\mathcal{G}_k = (V, E_k) = (V, \{(x, y) \mid x, y \in V, i_x(y) \leq k\}),$$

where $i_x(y)$ is the index of point y in a list of points $V \setminus x$ sorted in increasing order of distance from x . This method is simple to implement, but does display a greater degree of sensitivity to variations in data point sampling density.

Once the distance-weighted nearest neighbour graph has been constructed, using either the ε -Isomap or k -Isomap method, distances between arbitrary data points, $d_G(x, y)$, are defined by shortest paths in the graph. These shortest paths can be determined using standard graph algorithms; here, we use Floyd's all-sources shortest paths algorithm (Aho et al., 1983). Although this algorithm has time complexity $O(n^3)$, it is good enough for our purposes since the number of data points is not large ($n \leq 6000$). More efficient algorithms, for instance a Fibonacci heap-based implementation of Dijkstra's algorithm, give better performance for larger datasets. Asymptotic convergence results exist showing that the difference between the approximation $d_G(x, y)$ and the true geodesic distance, $d_M(x, y)$, tends to zero in a probabilistic sense as the density of data points increases (Bernstein et al. 2000¹). From these results, one can derive a required data point density to achieve any desired accuracy for $d_G(x, y)$. Unfortunately, these results are of limited use in practice. One usually starts with a set of data with a given, probably inhomogeneous, sampling density, and one would like to choose k or ε so as to produce robust results from Isomap. This is difficult, and the best approach seems to be a brute force sensitivity analysis over reasonable ranges of k and/or ε to probe different scales in the data.

4.1.2 Multidimensional scaling

Once the approximate geodesic distance function $d_G(x, y)$ has been found, a multidimensional scaling (MDS) proce-

¹Bernstein, M., de Silva, V., Langford, J. C., and Tenenbaum, J. B.: Graph approximations to geodesics on embedded manifolds, <http://isomap.stanford.edu/BdSLT.pdf>, pre-print, 2000.

ture is applied. This procedure results in an eigenvalue spectrum that can be examined to determine the dimensionality of the data manifold. It also calculates embeddings of the data points into low-dimensional Euclidean spaces.

MDS (Borg and Groenen, 1997) is a statistical technique that takes as input distance or dissimilarity measures for a set of data points and attempts to find points in Euclidean space such that the Euclidean distances between the output points correspond to the distance or dissimilarity values between the input points. Both PCA and Isomap can be considered within this framework. For PCA, the input distances are Euclidean distances in the input data, so that MDS leads to an orthogonal transformation of the data. For an idealisation of Isomap where the input distances are exact geodesic distances in the data manifold, MDS leads to an isometric transformation of the data.

The form of MDS used in Isomap is usually referred to as *classical scaling* (Torgerson, 1952; Gower, 1966; Borg and Groenen, 1997). As input, we require a distance or dissimilarity measure $d_{ij}=d(\mathbf{x}_i, \mathbf{x}_j)$ calculated between the n data points, $\mathbf{x}_i \in \mathbb{R}^m$. The distance function must satisfy the usual conditions for distances: $d_{ii}=0, d_{ij}=d_{ji}, d_{ik} \leq d_{ij}+d_{jk}$.

From the distance function, we form a matrix of squared distances $(\mathbf{D}^{(2)})_{ij}=d_{ij}^2$. To this matrix we then apply a double centring transformation, using a centring operator $\mathbf{J}=\mathbf{I}-n^{-1}\mathbf{1}\mathbf{1}^T$, with \mathbf{I} being the $n \times n$ identity matrix and $\mathbf{1}$ an n element vector of ones. The centring transformation is

$$\mathbf{Z}^{(2)} = -\frac{1}{2}\mathbf{J}\mathbf{D}^{(2)}\mathbf{J}.$$

A simple calculation shows that, if d_{ij} is a Euclidean distance function, then $\mathbf{Z}^{(2)}$ is the matrix of scalar products between the vectors $\{\mathbf{x}_i\}$, i.e. $(\mathbf{Z}^{(2)})_{ij}=\mathbf{x}_i \cdot \mathbf{x}_j$. For centred data, i.e. data for which the mean of the \mathbf{x}_i is zero, $\mathbf{Z}^{(2)}$ then corresponds to the covariance matrix normally used for PCA. For non-Euclidean distance functions, the matrix $\mathbf{Z}^{(2)}$ encodes comparable information about the distribution of distances between the data points.

Next, the eigenvector decomposition of the scalar product matrix $\mathbf{Z}^{(2)}$ is calculated, as $\mathbf{Z}^{(2)}=\mathbf{Q}\mathbf{\Lambda}\mathbf{Q}^T$, where $\mathbf{\Lambda}=\text{diag}(\lambda_1, \dots, \lambda_n)$ is a diagonal matrix with the eigenvalues of $\mathbf{Z}^{(2)}$ along its leading diagonal, and \mathbf{Q} is a matrix with the eigenvectors of $\mathbf{Z}^{(2)}$ as its columns. The usual hope is that, if the eigenvalues λ_i are sorted in order of decreasing magnitude, $\lambda_p \gg \lambda_{p+1}$ for some $p < m$ and we can approximate the matrix $\mathbf{Z}^{(2)}$ by projection onto the subspace spanned by the p leading eigenvectors. If we denote the matrix of the first p eigenvalues by $\mathbf{\Lambda}_+$ and the first p columns of \mathbf{Q} by \mathbf{Q}_+ , then the matrix of p -dimensional reduced coordinates for the data points is given by $\mathbf{X}=\mathbf{Q}_+\mathbf{\Lambda}_+^{1/2}$. Equivalently, denoting the eigenvectors of $\mathbf{Z}^{(2)}$ by \mathbf{q}_k , the k th coordinate of the i th data point in a p -dimensional reduced representation is

$$y_i^k = \sqrt{\lambda_k}q_k^i, \quad k = 1, \dots, p. \tag{1}$$




Fig. 5. The Swiss roll dataset.

This procedure is essentially that followed in PCA, apart from possible differences in data normalisation, but there are two problems, one common to all MDS algorithms and one important only in the more general setting relevant to Isomap. First, there is no guarantee that there is a gap in the eigenvalue spectrum of $\mathbf{Z}^{(2)}$, making it difficult to decide on a reduced dimensionality for the data. Second, the procedure described here is dependent on the non-negativity of the eigenvalues of the matrix $\mathbf{Z}^{(2)}$. In the case of PCA, positive semi-definiteness of $\mathbf{Z}^{(2)}$ is guaranteed by the use of Euclidean distances between data points, but in the more general case of Isomap, this is no longer the case. For an exact calculation of geodesic distances in an intrinsically flat manifold, the distance metric is Euclidean and $\mathbf{Z}^{(2)}$ is positive semi-definite. In Isomap, geodesics are calculated only approximately, and errors associated with the approximation are often enough to render $\mathbf{Z}^{(2)}$ non-positive semi-definite, yielding negative eigenvalues in the MDS procedure. Another possible source of negative eigenvalues in Isomap is the structure of the data manifold. Isomap assumes that the data manifold is globally isometric to an open, connected, convex subset of Euclidean space (Donoho and Grimes, 2005). Data manifolds that are not convex (i.e. that do not contain all geodesics connecting points lying in the manifold – an example is a two-dimensional surface with a hole, which is then not simply connected) or that possess non-zero intrinsic curvature do not satisfy these assumptions and have geodesic distance functions that lead to $\mathbf{Z}^{(2)}$ matrices with negative eigenvalues.

Eigenvalues in MDS and, in particular, in PCA, are customarily interpreted as the proportion of the total data variance explained by a particular mode. Clearly, negative eigenvalues cannot be interpreted as variances. One approach is to ignore any negative eigenvalues, assuming them to be the result of noise in the data or errors in the geodesic distance approximation. A more satisfactory approach is to observe that negative eigenvalues are always small and always paired with positive eigenvalues of similar magnitude, constituting the tail of the eigenvalue distribution. The presence of negative eigenvalues can still be considered a form of noise, but

the position in the eigenvalue spectrum of the first negative eigenvalue can be used as a cut-off point for considering the reduced dimensionality of the data. According to this view, no positive eigenvalue appearing after a negative eigenvalue can correspond to a real dimension in the reduced dimensionality data. The justification for this interpretation is simply that negative eigenvalues cannot be interpreted as variances, cannot be used in Eq. (1) to calculate reduced coordinates and so must be neglected. Some complication is entailed by this viewpoint, since it is no longer possible to use a simple measure of explained variance such as $c_p = \sum_{i=1}^p \lambda_i / \text{Tr } \Lambda$ because the trace of the eigenvalue matrix no longer measures the total variance in the data due to the presence of the negative eigenvalues. It is thus not possible to use an explained variance threshold to infer the dimensionality of the data and to choose a set of modes on to which to project. Here, we use a different approach, finding a pair of straight lines with a “knee” that best fits the MDS eigenvalue spectrum in a least squares sense and taking the dimensionality of the data to lie at the knee. This approach, which is easy to understand and proves to be reasonably robust, is explained in detail in Sect. 4.2.

4.1.3 Computational complexity

The two main computational bottlenecks in the Isomap algorithm are the computation of the nearest neighbour graph and the final MDS eigenvalue problem, which, for n data points, involves finding the leading eigenvalues and eigenvectors of an $n \times n$ matrix. A naive implementation using a dense eigenvalue solver has a computational cost that scales as $O(n^3)$. Here, we have datasets with $n \leq 6000$, and use the Anasazi iterative eigenvalue solver from the Trilinos project (Baker et al., 2008², Heroux et al., 2005). The block Krylov-Schur scheme in Anasazi is able to find the first fifteen eigenvalues and eigenvectors of a 6000×6000 matrix in a time entirely negligible compared to the time required for the all-sources shortest paths calculation used to approximate geodesic distances in the data manifold. For still larger problems, an adaptation of Isomap exists using a smaller number of *landmark* points (de Silva and Tenenbaum, 2002), but this refinement did not prove necessary here.

4.2 Isomap sensitivity

The Isomap algorithm has a single tunable parameter, the number of nearest neighbours used to construct the graph on which the approximate geodesic calculation is based. A natural issue to investigate is to what extent results inferred from Isomap depend on this parameter.

To explore some of the implications of sensitivity to this parameter choice, we use a simple “Swiss roll” dataset,

²Baker, C. G., Hetmaniuk, U. L., Lehoucq, R. B., and Thornquist, H. K.: Anasazi software for the numerical solution of large-scale eigenvalue problems, ACM T. Math. Software, in press, 2008.

representing a two-dimensional manifold embedded in \mathbb{R}^3 . Figure 5 illustrates the essential features of this data – the manifold in which the data points lie is intrinsically flat, but curled up so that points far apart according to the intrinsic geodesic metric in the manifold are close together as measured by the Euclidean metric in the embedding space. The implications of this for the construction of the Isomap nearest neighbour graph are clear: choosing too large a number of nearest neighbours k or too large a radius ε will cause points on adjacent but separate leaves of the manifold to be identified as nearest neighbours, leading to an incorrect identification of the topology of the data manifold.

Figure 6 shows results from Isomap sensitivity studies using the Swiss roll data, one for ε -Isomap (Fig. 6a) and one for k -Isomap (Fig. 6b). Each plot shows MDS eigenvalue spectra in contour form, as a function of eigenvalue number and the nearest neighbour parameter (ε or k).

As previously mentioned, if negative eigenvalues are present in the MDS spectrum, they must be excluded from any dimensionality reduction, since they cannot be viewed as measures of explained variance, and cannot be interpreted in terms of a lower-dimensional real manifold. The areas filled in grey in Fig. 6 indicate regions of eigenvalue space that are forbidden by this condition. No eigenvalues beyond the first negative eigenvalue can be part of a real lower-dimensional representation of the data. Given this constraint, the dimensionality of the data is estimated by looking for a “knee” in the eigenvalue spectrum, and is indicated in Fig. 6 by a thick red line.

In both plots in Fig. 6, there is a change in behaviour of the eigenvalue spectra as the nearest neighbour parameter is varied: at $\varepsilon \approx 3.6$ or $k=7$, there is a distinct step change in the spectra. For neighbourhood sizes below the threshold, the convergence of the eigenvalue spectra is quicker than for values above the threshold. Consequently, the dimensionality estimates inferred are lower for neighbourhood sizes below the threshold. For the ε -Isomap results, this effect reflects the fact that, in the norm used here, the separation between adjacent leaves of the spiral in the Swiss roll data is about 3.6. For neighbourhood radii smaller than this, the nearest neighbour connections in the distance-weighted graph used to approximate geodesics are confined to the surface of the manifold. For larger neighbourhood radii, the neighbourhoods spill over between adjacent leaves of the manifold. Varying the neighbourhood parameter probes different scales in the data. Smaller values of ε pick out smaller scale structures and detect the separation between the leaves of the manifold. Larger values of ε do not resolve this fine structure and see the data as an amorphous cloud of points. Small values of ε thus give $p=2$, the true dimensionality of the embedded manifold, while larger values give $p=3$, the dimension of the embedding space.

Similar conclusions can be drawn from the k -Isomap results (Fig. 6b), though here the value of k at which the transition from $p=2$ to $p=3$ occurs is harder to interpret. The

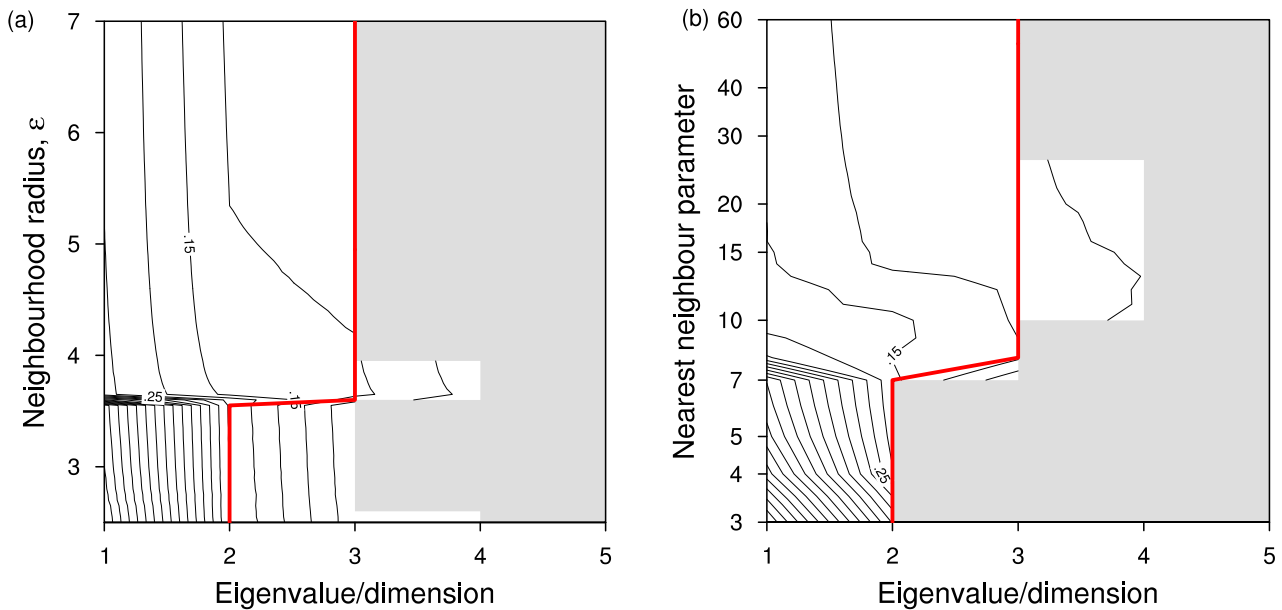


Fig. 6. Isomap eigenvalue convergence and dimension estimates for the Swiss roll dataset, a two-dimensional manifold embedded in \mathbb{R}^3 . Black contours show MDS eigenvalue spectra normalised by the overall largest eigenvalue, as a function of eigenvalue number and neighbourhood radius ε (a) or nearest neighbour count k (b) (logarithmic axis). The grey areas indicate regions of the eigenvalue spectra that are not useful for dimensionality reduction because of the presence of negative eigenvalues. The thick red line shows the data dimensionality, estimated from the eigenvalue spectra as described in the main text. The “true” dimensionality of the dataset is two.

transitional value $k=7$ is the number of neighbours, on average, that a data point has within a radius of $\varepsilon \approx 3.6$, but this number is subject to large sampling variability, giving a slightly rougher transition for k -Isomap than ε -Isomap. The dataset used here has 1000 points, chosen to be comparable in size to the equatorial Pacific SST time series examined below, and this relatively small number of points in \mathbb{R}^3 leads to a wide range of variability in the distance from a point to its nearest neighbour (~ 0.02 – 2.13). There is thus a range of values of k for which the k nearest neighbours of some points all lie on the same leaf of the manifold while the k nearest neighbours of other points span more than one leaf. Despite this, the dimensionality estimates are the same as for ε -Isomap, i.e. $p=2$ for $k \leq 7$ and $p=3$ for $k > 7$.

It should be noted that the dimensionality inferred from Isomap depends to a certain extent on subjective factors. Although there is no need to choose a total cumulative explained variance to select the number of leading eigenvectors to consider, as is sometimes done with PCA, the condition for locating a “knee” in the eigenvalue spectrum is quite delicate. Here, we approximate the spectrum with a pair of lines with a kink at a selected eigenvalue, then choose the knee to be at that point whose fitted lines give the smallest RMS error when points on the lines are compared to the true eigenvalues. This approach substantially follows recommendations in Borg and Groenen (1997), but there are other methods that could equally be used.

The main conclusion to draw from this is that, at least in the case of the simple dataset used here, Isomap can probe the dimensionality of a lower-dimensional dataset embedded nonlinearly in a higher-dimensional space quite well. In this case, there is relatively little dependence of the results on the nearest neighbour parameter ε or k and what dependence is seen is well understood in terms of known characteristics of the dataset. The changes in MDS eigenvalue spectra seen as one varies the nearest neighbour neighbourhood size indicate how the method is probing the dataset at different scales. This dependence on the parameter ε or k can be viewed as a disadvantage (some value of k or ε needs to be chosen and there is no clear a priori method to do this) or an advantage (by varying k or ε , we can probe different scales to get a better idea of the underlying structure of our data). The results from ε -Isomap are easier to interpret because of the propensity for k -Isomap results to be influenced by data sampling variability, although k -Isomap is easier to use since there is no need to determine a suitable range for ε . The main impediment to performing the type of sensitivity analysis illustrated here is computing resource, since Isomap decompositions of the data for a large number of neighbourhood sizes are needed to form a clear picture of the structure of the variation in results with neighbourhood size.

In the sections below showing Isomap results for Pacific SST time series, sensitivity results are presented in parallel with other Isomap results to give some feeling for the

Table 2. Isomap dimensionality estimates for tropical Pacific SST data, for raw SSTs and SST anomalies. Values shown are the smallest and largest dimensionalities recovered by examining the Isomap eigenvalue spectra as the neighbourhood size parameter k or ε is varied.

| Dataset | Raw SST | | SST anomaly | |
|------------------|---------------|-----|---------------|-----|
| | ε | k | ε | k |
| Observations | 2–4 | 2–4 | 2–3 | 1–2 |
| BCCR-BCM2.0 | 2–3 | 2–3 | 2–4 | 2–2 |
| CCSM3 | 1–3 | 1–3 | 2–4 | 2–2 |
| CGCM3.1(T47) | 1–2 | 1–3 | 1–4 | 1–1 |
| CGCM3.1(T63) | 2–2 | 1–2 | 1–4 | 1–1 |
| CNRM-CM3 | 2–4 | 2–5 | 2–4 | 2–2 |
| CSIRO-Mk3.0 | 1–3 | 1–4 | 1–4 | 2–2 |
| ECHO-G | 2–4 | 4–5 | 1–4 | 2–2 |
| FGOALS-g1.0 | 3–4 | 1–4 | 2–3 | 2–5 |
| GFDL-CM2.0 | 2–3 | 2–3 | 1–1 | 1–2 |
| GFDL-CM2.1 | 2–3 | 2–4 | 1–2 | 1–2 |
| GISS-EH | 2–3 | 1–3 | 1–4 | 1–2 |
| INM-CM3.0 | 2–3 | 2–3 | 1–4 | 2–2 |
| IPSL-CM4 | 2–2 | 1–3 | 2–4 | 2–2 |
| MIROC3.2(hires) | 2–2 | 1–2 | 1–4 | 2–2 |
| MIROC3.2(medres) | 2–3 | 1–3 | 1–3 | 1–2 |
| MRI-CGCM2.3.2 | 2–4 | 3–4 | 1–2 | 1–2 |
| UKMO-HadCM3 | 2–4 | 3–5 | 2–3 | 2–2 |
| UKMO-HadGEM1 | 2–3 | 2–3 | 1–4 | 1–2 |

robustness of the method and the variability of the results with respect to the neighbourhood size. In general, the results are more dependent on neighbourhood size for the more complex tropical Pacific SST data, and the corresponding dimensionality estimates are less certain.

5 Results and discussion

All of the results reported here are based on the use of the full length of the model SST time series available, as listed in Table 1. Isomap eigenvalue spectra were also calculated for sub-segments of each dataset, consisting of 50, 25 and 10 year segments of the total available data, in order to determine the sensitivity of our results to time series length. The results (data not shown) indicate that there is little variation in the Isomap eigenvalue spectra we calculate, at least for 50 or 25 year sub-segments, leading us to conclude that our results are reasonably robust with respect to variations in the amount of data available.

5.1 Analysis for raw SSTs

In this section, we present Isomap results for tropical Pacific SSTs from observational and model datasets. In performing PCA, it is common to use SST anomalies, so removing the influence of the annual cycle. Isomap results for SST anomalies

are presented in Sect. 5.2, allowing for direct comparison between PCA and Isomap, but here, one of the things we wish to explore is the extent to which Isomap is able to determine the coupling between ENSO and annual variability in the tropical Pacific. This coupling is one factor lost in the customary anomaly-based PCA approach.

In this section we use SSTs and in the next, SST anomalies, from the region 125° W–65° W, 20° S–20° N, normalising each dataset to zero mean and unit standard deviation at each spatial point. This choice of normalisation is used throughout to permit direct comparison with the earlier work of Gámez et al. (2004).

The leading modes of variability in tropical Pacific SSTs are the annual cycle and ENSO, and we expect Isomap to pick these out. As in the case of the Swiss roll data, it is useful to examine the sensitivity of Isomap results to variations in the ε or k neighbourhood size parameters. Figure 7 displays Isomap sensitivity plots for observational SST data (Figs. 7a and 7d) and two selected models (Figs. 7b, 7c, 7e and 7f). Compared to the Swiss roll results (Fig. 6), the eigenvalue spectra and corresponding dimensionality estimates for the SST data show more variation with Isomap neighbourhood size. The ranges of k and ε used in Fig. 7 are selected to correspond as far as possible, but it is difficult to relate results for any particular value of k to those for any particular value of ε or vice versa because of the variability in distances between data points. One common feature in the ε -Isomap plots in Fig. 7 is that the regions of negative eigenvalues in the Isomap spectra disappear as the neighbourhood size increases. This reflects the equivalence of Isomap with a large neighbourhood size to PCA under suitable conditions of data normalisation: in the limit of infinite neighbourhood size, the geodesic distance approximation used in Isomap collapses to the use of the original Euclidean distances between the data points, so is equivalent to PCA. The same effect would also be seen in the k -Isomap results for $k \approx 1000$, the number of data points used.

Despite the high embedding dimension of the data (essentially the number of non-land points in the study region, m in Table 1), the dimensionality estimates inferred from Isomap in Fig. 7 are rather low. This is true for all models examined and for the observational data. Table 2 shows the range of dimensionality estimates inferred for each dataset. For raw SSTs, across all datasets the dimensionality estimates range from 1 to a maximum of about 5. The eigenvalue spectra here converge rapidly because the leading modes of variability are overwhelmingly larger in amplitude than the other modes. The coherent variation of SST patterns in the tropical Pacific can easily be represented by a small set of modes. The convergence of the Isomap eigenvalue spectra is rather quicker than the convergence of eigenvalue spectra for PCA performed in a comparable setting, i.e. using raw SST data rather than SST anomalies, as shown in Gámez et al. (2004). This quicker convergence can be ascribed to better representation of the nonlinear ENSO variability by Isomap than by

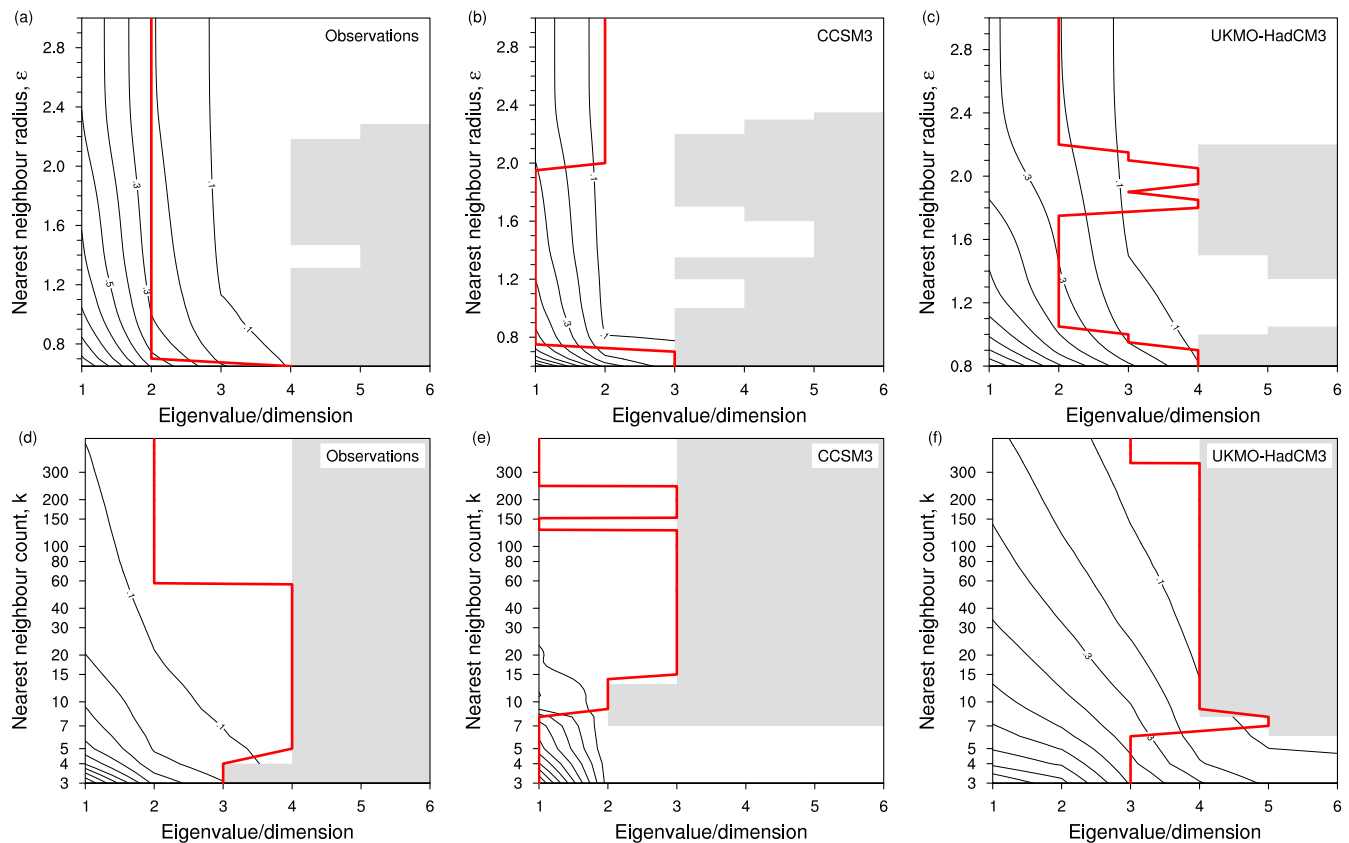


Fig. 7. Isomap eigenvalue convergence and dimension estimates for tropical Pacific raw SSTs, from observations (**a** and **d**), and two models, CCSM3 (**b** and **e**) and UKMO-HadCM3 (**c** and **f**). Black contours show MDS eigenvalue spectra normalised by the overall largest eigenvalue, as a function of eigenvalue number and neighbourhood radius ε (a–c) or nearest neighbour count k (d–f) (logarithmic axis). The grey areas indicate regions of the eigenvalue spectra that are not useful for dimensionality reduction because of the presence of negative eigenvalues. The thick red line shows the data dimensionality, estimated from the eigenvalue spectra.

PCA. The PC scatter plots shown earlier (Fig. 4) demonstrate that ENSO variability is probably not a linear Gaussian phenomenon, so this is expected.

The range of dimension estimates shown in Table 2 for SST observations (2–4) is what we would expect, including two dimensions to describe the periodic annual cycle and one or two for ENSO variability. Here, two degrees of freedom are expected for the annual cycle because of the geometry of manifolds that can be faithfully represented by Isomap. The globally isometric transformation used by Isomap permits it to represent only simple Euclidean coordinates and not periodic coordinates, meaning that any periodic phenomenon requires at least two degrees of freedom. There is no equivalent to the “circular” bottleneck layer NLPCA procedure described in Hsieh (2004) that allows periodic coordinates to be extracted directly. For ENSO variability, as well as the leading degree of freedom usually represented by the NINO3 SST index, previous studies have identified a second degree of freedom, varying in quadrature with the first, correspond-

ing to changes in the equatorial Pacific warm water volume (Burgers, 1999; Kessler, 2002; McPhaden, 2003).

Some model results show lower dimensional behaviour than this, including CCSM3 and CGCM3.1 (both T47 and T63). In the case of CCSM3, the reason for this behaviour is seen in the NINO3 power spectra in Fig. 2. Here, the observational data show a broad peak in the ENSO power band (2–7 years). CCSM3, however, has a sharper peak at almost exactly 2 years, displaying a mode of variability rather different from observed ENSO variability. In the Isomap analysis, this biannual variability is aliased with the annual cycle, and no distinct ENSO variability is detected. The situation with the CGCM3.1 models is different. Here, the NINO3 power spectrum shows essentially no peak in the ENSO frequency band. It is not clear what is happening here, but it may be relevant that the equatorial SST climatology in both CGCM3.1 models is poor, showing little or no gradient across the Pacific basin (Fig. 1).

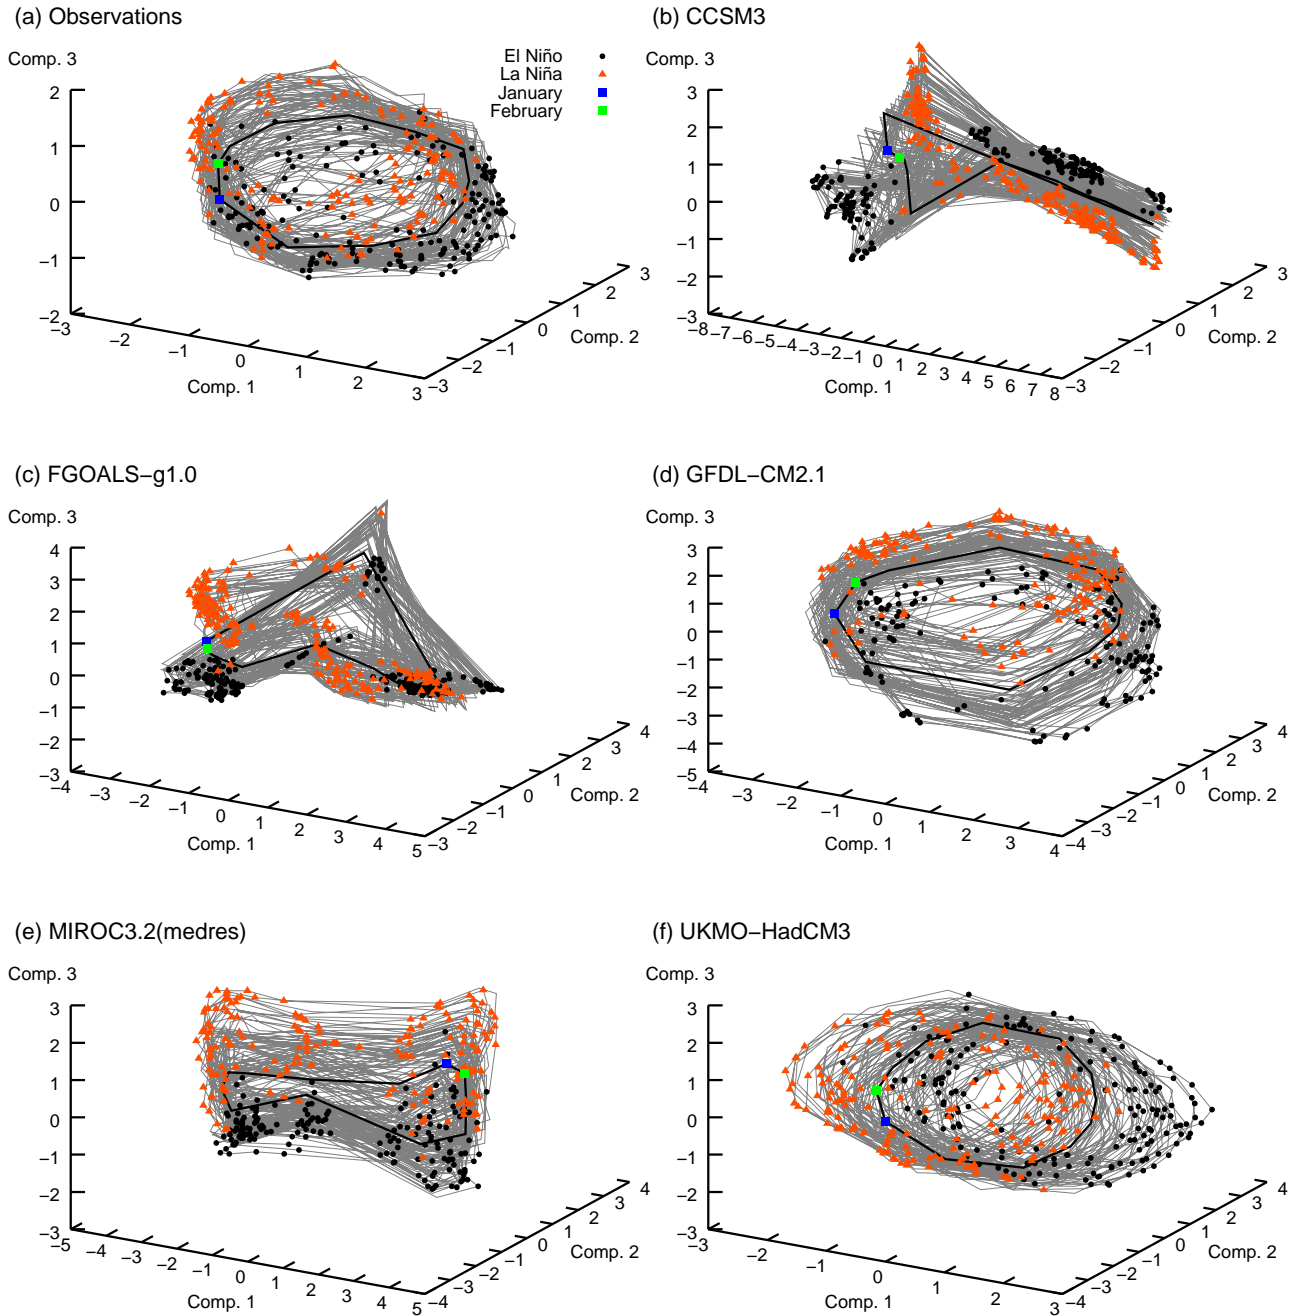


Fig. 8. Three-dimensional embeddings of Isomap raw SST results for observations (a) and selected models (b–f). Light grey lines join data points representing adjacent months in the SST time series. The mean annual cycle is shown as a thicker line with January and February highlighted in blue and green respectively. Points are identified as El Niño (black dots) or La Niña (red triangles) events based on the corresponding NINO3 SST index time series for each dataset. For clarity, only 100 years of data is plotted here for each model.

Table 3. Correlation coefficients between NINO3 SST index and warm water volume (WWV) and Isomap components from k -Isomap with $k=7$: for raw SSTs, the correlation between rotated Isomap component #3 and NINO3 and between rotated Isomap component #4 and WWV; for SST anomalies, the correlation between Isomap component #1 and NINO3 and between Isomap component #2 and WWV. Blank entries occur where the Isomap eigenvalue spectrum in a particular case does not have enough positive eigenvalues to form an embedding of the required dimensionality.

| Dataset | Correlation | | | |
|------------------|-------------|-------|-------------|------------------|
| | Raw SST | | SST anomaly | |
| | NINO3 | WWV | NINO3 | WWV |
| Observations | 0.822 | 0.031 | 0.841 | 0.242 |
| BCCR-BCM2.0 | 0.835 | | 0.820 | 0.153 |
| CCSM3 | 0.047 | 0.284 | 0.901 | 0.245 |
| CGCM3.1(T47) | 0.223 | | 0.824 | 0.021 |
| CGCM3.1(T63) | 0.225 | 0.148 | 0.814 | 0.284 |
| CNRM-CM3 | 0.824 | 0.225 | 0.776 | 0.228 |
| CSIRO-Mk3.0 | 0.746 | 0.227 | 0.717 | 0.407 |
| ECHO-G | 0.907 | 0.681 | 0.935 | 0.646 |
| FGOALS-g1.0 | 0.793 | | 0.776 | 0.430 |
| GFDL-CM2.0 | | | 0.857 | 0.435 |
| GFDL-CM2.1 | 0.859 | | 0.853 | 0.665 |
| GISS-EH | 0.652 | | 0.665 | 0.116 |
| INM-CM3.0 | 0.730 | | 0.744 | 0.446 |
| IPSL-CM4 | 0.844 | | 0.852 | n/a ^a |
| MIROC3.2(hires) | 0.236 | 0.171 | 0.686 | 0.070 |
| MIROC3.2(medres) | 0.646 | | 0.785 | 0.087 |
| MRI-CGCM2.3.2 | 0.861 | | 0.909 | 0.511 |
| UKMO-HadCM3 | 0.804 | 0.093 | 0.809 | 0.012 |
| UKMO-HadGEM1 | 0.747 | | 0.752 | 0.274 |

^a Ocean temperature data required to calculate warm water volume for IPSL-CM4 is not available.

Once we select a dimensionality for embedding of Isomap results, we can calculate reduced coordinates using Eq. (1). Here, we initially select an embedding dimensionality of three, both because this lies in the range derived from the Isomap eigenvalue spectra and because it is the highest dimensionality of data we can easily visualise. Figure 8 illustrates three-dimensional embeddings for SST observations and a selection of models. The Isomap results shown are all for k -Isomap with $k=7$. The plots show the data as a time series, with points adjacent in time connected by thin grey lines. The mean annual cycle is shown as a thicker line with January and February highlighted for orientation. Points identified as El Niño or La Niña events on the basis of the NINO3 SST index are picked out in colour. For clarity, only 100 years of the Isomap results are plotted in each case. Concentrating on the observations first, it can be seen that Isomap correctly identifies the annual cycle (represented by motion about the roughly cylindrical region occupied by the data points) and at least one other form of variability (repre-




Fig. 9. Time series of NINO3 SST index (black) and rotated Isomap component #3 (red) for observations (a) and selected models (b–e). An arbitrary 20 year slice of data is shown in each case.

sented by motion approximately in the direction of the axis of the cylindrical region). The clustering of El Niño and La Niña points indicates that this second mode of variability corresponds to ENSO and generally lies along the direction orthogonal to the annual cycle in the embedding coordinates. Following Gámez et al. (2004), the role of the “axial” mode can be clarified by rotating the Isomap embedding to bring the mean annual cycle into the x - y coordinate plane. In this rotated coordinate system, variations in the z -direction record the “axial” variability in the original embedding coordinates (see Appendix for details of this rotation procedure). Time series plots of the rotated third Isomap component for observations and four of the models selected here are shown in Fig. 9. The rotated Isomap component #3 time series are plotted in parallel with time series of the NINO3 SST index, recording ENSO variability. For the observations, in Fig. 9a, it is clear that rotated Isomap component #3 quite accurately captures ENSO variability in the input SST data. In this case, Isomap has thus extracted the most important

modes of variability in tropical Pacific SSTs, the annual cycle and ENSO, starting from high-dimensional input data. We can also go further and attempt to extract the second degree of freedom in ENSO variability, usually identified with the equatorial Pacific WWV (Kessler, 2002; McPhaden, 2003), by examining a four-dimensional embedding of the Isomap results. The same sort of rotation procedure can be applied to remove the influence of the annual cycle variability on both Isomap components #3 and #4 (see Appendix for details). Correlation coefficients between Isomap rotated component #3 and the NINO3 SST index and between Isomap rotated component #4 and WWV are shown in Table 3. For the observational data, the NINO3 correlation is high, as would be expected from Fig. 9a, but the correlation between rotated Isomap component #4 and WWV is very low. It thus appears that rotated Isomap component #4 here does not capture this second degree of ENSO variability.

Although the fact that Isomap appears to capture the annual cycle variability and at least some aspects of ENSO variability is unsurprising, the data-driven nature of Isomap makes it useful for comparison of model results with observations and for inter-model comparison. We apply the same three-dimensional embedding to selected model results in Figs. 8b–f. Results for a number of the models shown (GFDL-CM2.1, MIROC3.2(medres) and UKMO-HadCM3) are similar to observations, with a clear three-dimensional structure to the data embedding, cleanly picking out the annual cycle and ENSO, with distinct clustering of El Niño and La Niña events. For the other two models illustrated, CCSM3 and FGOALS-g1.0, the three-dimensional Isomap embedding reveals data manifolds of significantly different form to that of the observations. As noted earlier, for CCSM3 this is due to excessively regular interannual variability in tropical Pacific SSTs that appears to be aliased with the annual cycle. For FGOALS-g1.0, the situation appears to be similar. The FGOALS-g1.0 NINO3 power spectrum in Fig. 2 exhibits a narrow peak at a period of around 3.5 years, rather than a broad peak stretching across the 2–7 year ENSO power band. This narrowband signal is again likely to result in lower-dimensional behaviour in the Isomap results.

Time series of rotated Isomap component #3 alongside the NINO3 SST index are plotted for a smaller selection of models in Figs. 9b–e. Two of these cases, GFDL-CM2.1 (Fig. 9c) and UKMO-HadCM3 (Fig. 9e), are models whose three-dimensional Isomap embeddings show similar structure to observations. This is reflected in the rotated Isomap component #3 time series, which show good correlation with the NINO3 SST index. A good correlation is also seen for the results for FGOALS-g1.0 (Fig. 9b), despite the apparent degeneracy of the 3-D Isomap embedding in Fig. 8c. Despite the visual discrepancy between the FGOALS-g1.0 embedding results and the observations, it appears that the Isomap algorithm is still able to disentangle the annual and ENSO variability in the modelled SST data. The other model illustrated in Fig. 9 is MIROC3.2(medres) (Fig. 9d), which has

weaker ENSO variability, but still shows a reasonable correlation between rotated Isomap component #3 and the NINO3 SST index.

For models with strongly degenerate three-dimensional Isomap embeddings, such as CCSM3 (Fig. 8b) and CGCM3.1(T47), CGCM3.1(T63) and MIROC3.2(hires) (not shown), the rotated Isomap component #3 time series show little coherent variability, and certainly none that correlates with ENSO variability. Correlation coefficients between rotated Isomap component #3 and the NINO3 SST index are shown in Table 3 for all models along with observational data. The models showing good correlations are those models for which the three-dimensional Isomap embedding displays a similar structure to the observations, i.e. for which Isomap successfully extracts the annual cycle and an “orthogonal” ENSO mode. As for the observations, we can also attempt to identify a second degree of freedom of ENSO variability by examining four-dimensional Isomap embeddings. One problem here is that, for some of the models, the Isomap eigenvalue spectra do not have enough positive leading eigenvalues to provide a four-dimensional embedding – at least four positive leading eigenvalues are required. In the cases where a four-dimensional embedding of the Isomap results is possible, we conduct the same four-dimensional rotation as for the observations, to remove the annual variability from both rotated Isomap components #3 and #4, and then calculate correlation coefficients between the rotated Isomap components and the NINO3 SST index and simulated WWV time series, calculated as described in Sect. 2. As for the observations, the correlations between rotated Isomap component #4 and WWV for the models are generally rather low.

As noted at the beginning of this section, one reason for applying Isomap to raw SST data, as opposed to SST anomalies, was to determine the extent to which Isomap is able to identify the coupling between annual and ENSO variability in the tropical Pacific. Other, more direct, analyses of ENSO/annual cycle interactions reveal a strong influence of the magnitude of the annual cycle in the equatorial Pacific on ENSO variability (Guilyardi, 2006). On the basis of the results presented here, it appears that our Isomap analysis does not provide very much insight into this question.

5.2 Analysis for SST anomalies

In climatological contexts, PCA is normally applied to climate anomalies, i.e. to data from which the mean annual cycle has been removed. This was the case for the equatorial Pacific SST EOFs shown in Sect. 3. We can also apply Isomap to SST anomalies, thus providing results that are more directly comparable with the results of PCA than the raw SST Isomap analysis presented in the previous section. These results may also be slightly easier to interpret because rotation to remove the influence of the annual cycle is not required.

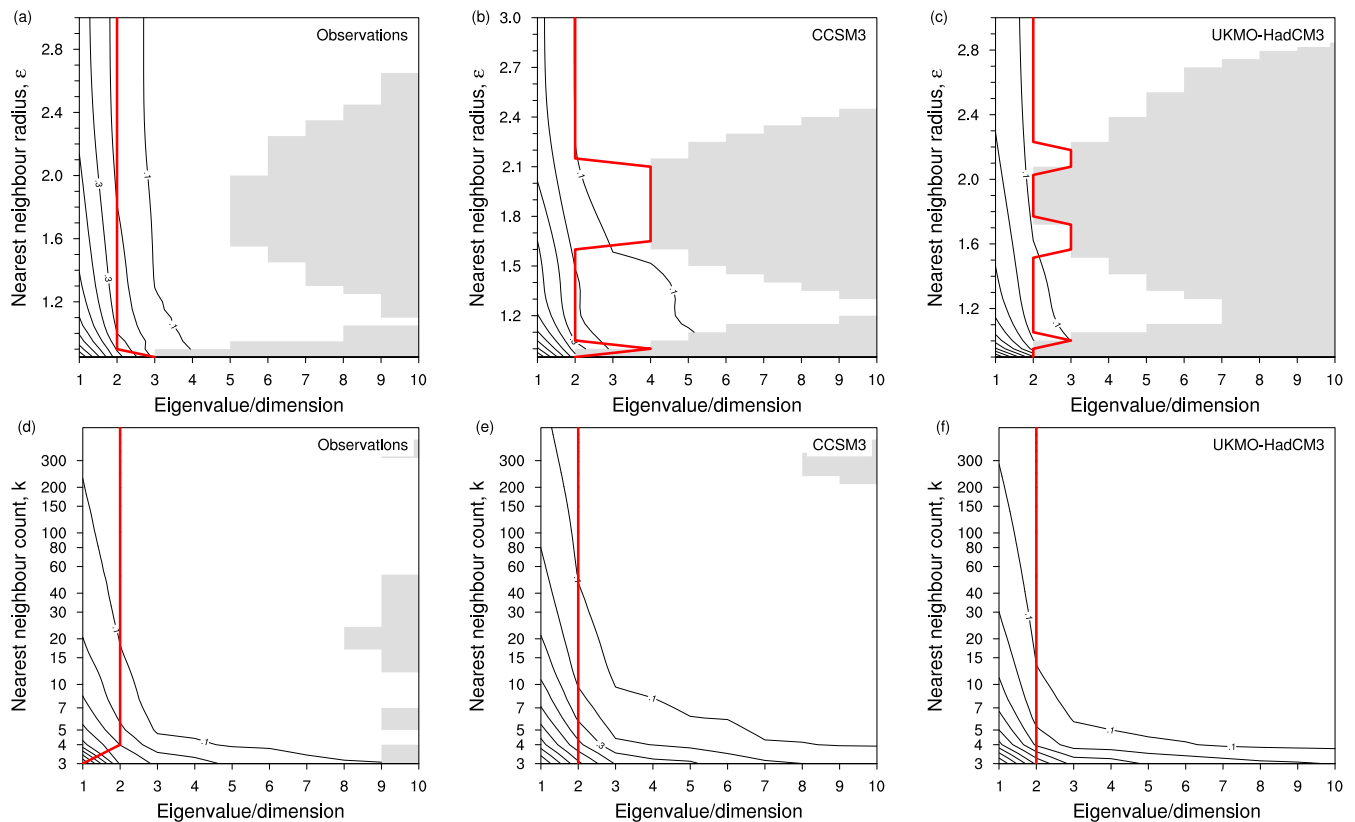


Fig. 10. Isomap eigenvalue convergence and dimension estimates for tropical Pacific SST anomalies, from observations (**a** and **d**), and two models: CCSM3 (**b** and **e**) and UKMO-HadCM3 (**c** and **f**). Black contours show MDS eigenvalue spectra normalised by the overall largest eigenvalue, as a function of eigenvalue number and neighbourhood radius ε (a–c) or nearest neighbour count k (d–f) (logarithmic axis). The grey areas indicate regions of the eigenvalue spectra that are not useful for dimensionality reduction because of the presence of negative eigenvalues. The thick red line shows the data dimensionality, estimated from the eigenvalue spectra.

As for the raw SST Isomap results, the sensitivity of the SST anomaly Isomap results to variations in the ε or k parameters can be examined. Results for observations and selected models are plotted in Fig. 10 and minimum and maximum dimensionality estimates derived from these plots are shown in Table 2. It can be seen that the dimensionality estimates for the SST anomaly data are all rather low, with only one model (FGOALS-g1.0) having a maximum dimensionality greater than two. This indicates that only one- or two-dimensional embeddings of the Isomap results are possible here.

We thus examine one-dimensional and where available, two-dimensional, embeddings of the Isomap results. The justification for this is that we expect ENSO to be the major mode of variability in the SST anomalies, with the first component of any embedding corresponding to the NINO3 SST index variability, and the second component to the warm water volume variation – looking at one- or two-dimensional embeddings should pick these features out. Table 3 shows correlation coefficients between SST anomaly Isomap components #1 and #2 and the NINO3 SST index and WWV time

series respectively. The strong correlations between Isomap component #1 and the NINO3 index here indicate that the one-dimensional Isomap embedding does a good job of identifying the leading mode of ENSO variability, where it exists. For most models, the degree of correlation between the SST anomaly Isomap component #1 and the NINO3 SST index is similar to the degree of correlation between the raw SST rotated Isomap component #3 and the NINO3 SST index. For a small number of models though (primarily CCSM3, but also CGCM3.1(T47), CGCM3.1(T63) and to a lesser extent, MIROC3.2(hires)), the correlation for the SST anomaly Isomap component #1 is much higher than for the raw SST results. A reasonable explanation for this phenomenon in the case of CCSM3 is that the ENSO signal in this model is very regular, with a periodicity of almost exactly two years (Fig. 2), so is likely to be strongly aliased with the annual cycle in the raw SST results. Removing the annual cycle and working with SST anomalies may lift this degeneracy, allowing the “true” ENSO signal to be detected, leading to a stronger correlation. This aliasing-based explanation is less

applicable to the other models displaying large differences between the raw SST and SST anomaly Isomap correlation coefficients, since they do not have the same sort of very regular ENSO variability as CCSM3.

As for the raw SST results, we can also seek a second degree of freedom of ENSO variability by examining correlations between the SST anomaly Isomap component #2 and the warm water volume time series from the models. Here, some of the correlations between the Isomap SST anomaly component #2 and WWV are somewhat better than in the raw SST case, but there is still large variability in the correlations, and there is no clear link between a high correlation and a “good” ENSO. For instance, the CMIP3 models identified by van Oldenborgh et al. (2005) as having the best ENSO (GFDL-CM2.1, MIROC3.2(hires), MIROC3.2(medres), UKMO-HadCM3) have correlations ranging from 0.012 to 0.665. Again, it seems difficult to draw any clear conclusions from these results.

Scatter plots of the first two Isomap components display similar patterns to the principal component scatter plots of Fig. 4, which seems to indicate that the MDS eigenvectors produced by Isomap are nonlinearly related, just as are the EOFs produced by PCA. We believe that this may be a signal of intrinsic curvature in the data manifold. Isomap relies on isometric transformations of the data points and is therefore only able to represent embeddings of intrinsically flat manifolds. Attempting to project a manifold with non-zero intrinsic curvature to a lower-dimensional space by an isometric transformation necessarily leads to distortion of the relationships between points in the manifold.

5.3 Model WWV variability

It is quite possible that the low correlations seen between the WWV and both the rotated Isomap component #4 for the raw SST data and Isomap component #2 for the SST anomaly data stem not from problems with the Isomap analysis itself, but from deficiencies in the phasing of NINO3 SST index and WWV variations in the models. These variations can be characterised by examining plots of the NINO3 SST index time series versus the WWV time series. For the observations, there is a relatively well-defined phase relationship between variations in NINO3 SST and the warm water volume, particularly during El Niño events. This can be seen in Fig. 11 (after Fig. 2 of Kessler, 2002), where the monthly observed NINO3 SST index is plotted versus the observed warm water volume for the period 1980–1999. Large El Niño events, phase locked to occur in boreal winter, are clearly identified as loops in the plot, with large excursions to positive NINO3 SST index being associated with a corresponding coherent variation in WWV. Also visible is the “loitering” of the system during the recharge of the warm pool before the beginning of the next El Niño event, where predictability is generally lower (Kessler, 2002; McPhaden, 2003).

Equivalent plots for most of the models (not shown) show no such coherent variation between the NINO3 SST index and the warm water volume. Some models show clear El Niño events, often with some degree of locking to a consistent phase of the annual cycle, but concomitant variations in the warm water volume comparable to those seen in the observations are rare. The phenomenon of enhanced dwelling in the recharge phase is also not clearly seen in the majority of the models. The notable exception to this pattern is GFDL-CM2.1, which shows a pattern with a very strong resemblance to that of the observational data.

In the face of this lack of coherent NINO3/WWV variation, it seems unrealistic to expect Isomap to pick out any degree of freedom in ENSO variability in most of the models that displays any coherence with WWV variations. In this context, it is perhaps notable that the model with the greatest correlation coefficient between SST anomaly Isomap component #2 and WWV (0.665) is GFDL-CM2.1 (Table 3).

One possible reason for the lack of a clear relationship between model WWV and NINO3 SST index variations, pointed out by a reviewer, is our choice of method for computing the WWV. Model biases in mean equatorial thermocline mean that the 20°C isotherm may not be the best measure of thermocline depth, and that an index based on the location of the maximum vertical temperature gradient may provide a more consistent representation of the thermocline depth.

6 Conclusions

We have examined the applicability of Isomap to climate data analysis in the context of an inter-model comparison of ENSO variability. Our analysis indicates that Isomap is able to capture some of the low-dimensional dynamics of ENSO variability in the datasets that we have examined, picking out the gross features in the data. In some cases, notably for CCSM3, but also for CGCM3.1(T47), CGCM3.1(T47) and MIROC3.2(hires), examination of three-dimensional embeddings of the raw SST Isomap results, both visually (e.g. Fig. 8b) and via correlations between rotated Isomap component #3 and the NINO3 SST index (Table 3) reveals an anomalously low dimensionality of modelled ENSO variability, apparently caused by too regular interannual SST variability in the tropical Pacific, leading to aliasing of the ENSO signal to the annual cycle. Although this aspect of the models can be identified via other analyses, it is encouraging that Isomap is able to detect the anomalous behaviour without prompting. Less encouraging is the fact that Isomap is able to capture only these gross features of ENSO variability in the models. The Isomap results do not show much in the way of variations between models, at least not in an easily interpreted form. They also do not capture the sometimes significant differences between modelled and observed ENSO behaviour revealed by a simple comparison of

model and observational SST EOFs. Calculation of correlations between Isomap results and equatorial Pacific warm water volume (WWV) time series for both observations and model simulations do not reveal any strong relationship between the degrees of freedom found by Isomap and the second degree of freedom of ENSO variability that is generally believed to be represented by variations in WWV. Better results from this point of view might be found by performing an Isomap analysis directly on modelled thermocline depths, rather than simply trying to correlate WWV derived from the thermocline depths with SST Isomap results, although the unrealistic thermocline depth variation in some of the models is likely to make this difficult.

A more subtle illustration of differences between PCA and Isomap is presented by a comparison of the sensitivity of Isomap and conventional PCA to small changes in the structure of tropical Pacific SST variability around the shift in ENSO behaviour that occurred in the mid-1970s (Fedorov and Philander, 2000; McPhaden et al., 2006). If the observational dataset that we use is split into a pre-1976 component and a post-1976 component, differences relating to this change in ENSO behaviour are clearly apparent in SST EOFs, with a shift to stronger El Niño events. However, an Isomap analysis shows no significant differences in eigenvalue spectra between the pre-1976 and post-1976 data (data not shown). We speculate that this difference in sensitivity between PCA and Isomap is due to the fact that the orthogonal transformations associated with PCA, being more geometrically “rigid” than the isometric transformations of Isomap, are less able to conform to subtle changes in the data manifold, thus highlighting these relatively small differences.

Although in some senses Isomap is a rather blunt tool, it appears that it may be useful for exploratory data analysis, particularly if there is reason to believe that the data in question really is nonlinear and not too high-dimensional. In such cases, Isomap may serve a purpose alongside more conventional techniques.

There are four further issues with the Isomap algorithm that deserve comment, and that can provide a basis for comparison between Isomap and the NLPCA method (Hsieh, 2004). First is the question of the sensitivity of the results of the nonlinear dimensionality reduction techniques to choices of parameters in the algorithms used. For Isomap, this means variations in the k or ε neighbourhood size parameter. The possibility of varying this parameter can be viewed as an advantage, since it provides a mechanism to probe different length scales in the data in a way that has no analogue in PCA. How useful this is depends on the complexity of the dataset: for the simple Swiss roll data, a two-dimensional manifold embedded in \mathbb{R}^3 , variation in k or ε probes the structure of the data quite successfully. For the more complex ENSO datasets, it is not at all clear what sort of structures are being probed as the neighbourhood size is varied, and there is little consistency between the results from different models.

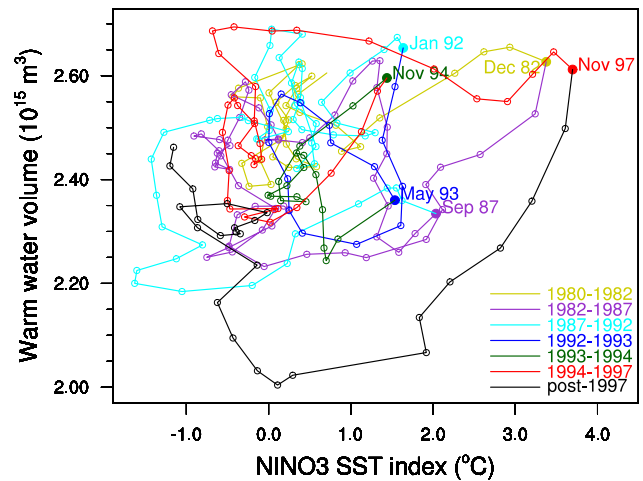


Fig. 11. Phasing of variations in observational NINO3 SST index (from the ERSST v2 dataset) and equatorial Pacific warm water volume (derived by Meinen and McPhaden, 2000 from sub-surface temperature data). Each point denotes a single month, ranging from January 1980 to December 1999. (After Fig. 2 of Kessler, 2002.)

In fact, from this point of view, the sensitivity of Isomap to the neighbourhood size is a clear disadvantage, since computational requirements generally restrict us to choosing a particular value of k or ε for our analyses, and there is no a priori reason to select one value over another. The situation for NLPCA is somewhat more complex than for Isomap, since there are a larger number of parameters involved: not only is there a choice of the exact structure of the network to be used (number of bottleneck nodes, special architectures for the bottleneck layer, number of nodes in hidden layers), but there are parameter choices involved in the protocol used to train the network without overfitting.

Second, results from Isomap are not easy to interpret if the underlying data manifold has a dimensionality higher than two or three. One example is an attempt to apply Isomap to mid-latitude tropospheric variability. Here, an Isomap analysis was performed for a monthly time series of 500 hPa geopotential height in the Atlantic sector. Isomap k/ε sensitivity studies (not shown) indicate a dimensionality of around 6 for the underlying data manifold. For manifolds of such high dimensionality, it is not possible to visualise the Isomap embeddings as we have done here for ENSO variability. Two- or three-dimensional projections are not sufficient to “unfold” the variability in the data, and the data points appear as an amorphous cloud of points. This situation also arises with PCA, if the eigenvalue spectrum converges slowly and many EOFs are required to explain a sufficient fraction of the data variance, but the linearity of PCA provides a partial solution. Linearity permits us to take single modes, EOFs, and treat them independently, one at a time. No such decomposition is possible for Isomap. This problem is not an inherent limitation of all nonlinear dimensionality reduction

techniques. For instance, NLPCA permits advance selection of the dimensionality to which the input data is to be reduced. Choosing a one-dimensional reduction gives the best nonlinear fit of a one-dimensional function to the input data, independent of the true dimensionality of the underlying data manifold, which the method makes no direct effort to ascertain. This approach allows for a “modal” analysis of the data, where nonlinear modes are stripped out of the input data one at a time. This type of analysis is not possible using Isomap because there is no way to control the dimensionality of the data reduction.

The third issue is shared with other nonlinear dimensionality reduction methods and is that it is generally difficult to produce plots showing spatial patterns of variability for nonlinear dimensionality reduction in the way that is done for PCA, where map plots of the leading EOFs are an important analytical tool. Such maps can be produced for PCA because real-valued EOFs essentially represent standing oscillations in the data, so a snapshot at any point in the oscillation from a positive pattern to a negative pattern records all the information about the spatial variability in the mode. For nonlinear methods, more general temporal variability is possible, and generally one needs to provide a set of spatial patterns corresponding to selected points on the reduced data manifold. Monahan (2001) and Hsieh (2004) demonstrate this approach for one-dimensional reduced manifolds, but for two-dimensional or larger manifolds, the number of spatial patterns needed becomes prohibitive.

The fourth point to note has been mentioned earlier when discussing the Isomap Pacific SST results, and this is the question of just what data manifolds a particular dimensionality reduction technique is capable of representing. As noted above, Isomap relies on a global isometric transformation of the original data space to derive a reduced Euclidean representation, meaning that only data manifolds that are globally isometric to Euclidean space can be faithfully represented by a reduced representation derived from Isomap. For NLPCA, the manifolds representable by the reduced representations depend on the structure of the bottleneck layer in the neural network. For a single bottleneck neuron, NLPCA can faithfully represent any open one-dimensional curve, for a “circular” bottleneck layer (two neurons, with values constrained to lie on the unit circle), closed one-dimensional curves can be represented faithfully, for two bottleneck neurons, general open two-dimensional surfaces can be represented, and so on. The complexity of interpreting the results of NLPCA increases quickly with the number of neurons in the bottleneck layer.

The essential problem with nonlinear methods such as Isomap is that there exist few theoretical results underpinning the numerical algorithms. For PCA, there are results identifying EOFs for at least some systems with normal modes of the system forced by random noise (North, 1984). These findings tie the numerical results of PCA directly to dynamical characteristics of the system under study. As far as we

know, there are no corresponding results for Isomap, or indeed any other nonlinear dimensionality reduction technique. There have been applications of Isomap to simple dynamical systems, where features observed in the Isomap results can be related to the dynamics of the system (Bollt, 2007), but no such studies exist for larger systems approaching the complexity of current climate models. Another approach to gaining analytical understanding is to explicitly construct data manifolds that can be exactly embedded by Isomap. Donoho and Grimes (2005) did this for an analytic representation of simple black-and-white images and developed several useful criteria for recognising classes of images whose data manifolds could be treated exactly by Isomap. It is not clear whether a similar approach to dimensionality reduction of dynamical systems would be fruitful.

Isomap is one of a large range of nonlinear dimensionality reduction techniques that have been developed in recent years (e.g., Roweis and Saul, 2000; Belkin and Niyogi, 2003; Donoho and Grimes, 2003; Weinberger and Saul, 2006; Lin et al., 2006). There is little a priori reason to choose one method over another for the analysis of any particular problem. For all of these methods, in order to assess which might be applicable in the analysis of climate data, it would be useful to have both the type of theoretical results mentioned above, and comparative studies applying the various methods to real problems. The results reported here are just one example, using one of the older, better-tested nonlinear dimensionality reduction techniques.

Appendix A Rotation of Isomap components

As described in Sect. 5.1, interpretation of three- and four-dimensional embeddings of raw SST Isomap results is clarified by rotating the components of the embeddings to separate the influence of annual variations (represented by rotated Isomap components #1 and #2) from the record of ENSO variability (as represented by rotated Isomap components #3 and #4). In this appendix, we explain the details of this rotation procedure, first for the three-dimensional case, then for the more complex four-dimensional case.

A1 Three-dimensional case

Consider a three-dimensional Isomap embedding of a monthly time series of N data items, resulting in a time series of 3-vectors \mathbf{y}_i , $i=1, \dots, N$, with components calculated from Eq. (1) of Sect. 4.1.2. Assuming that the time series covers a whole number of years, so that N is a multiple of 12, then the mean annual cycle for the embedding can be defined as $\bar{\mathbf{y}}_j$, $j=1, \dots, 12$, where

$$\bar{\mathbf{y}}_j = \frac{1}{N/12} \sum_{i=0}^{N/12-1} \mathbf{y}_{12i+j}.$$

In general, the points \bar{y}_j of the mean annual cycle will not lie in a single plane and, in particular, will not lie in the x - y coordinate plane. This means that each of the three components of the \bar{y}_j will vary over the course of the annual cycle, i.e. annual variability is “mixed into” each of the three components, even though only two Cartesian coordinates are strictly needed to represent the periodic annual variation.

In order to “unmix” the annual cycle from the third Isomap component, we may rotate the whole of the three-dimensional Isomap embedding to bring the mean annual cycle into the x - y coordinate plane, the hope then being that variations in the rotated Isomap component orthogonal to the x - y plane, i.e. variations along the z -axis, will represent interannual variability, specifically ENSO variability. As noted above, the mean annual cycle points \bar{y}_j do not generally lie in a plane, but we may identify a best-fit plane in a least-squares sense, and rotate this into the x - y plane. Although not perfect, this will lead to the most effective unmixing of annual variability from Isomap component #3. Figure A1 provides a schematic illustration of some of the details of the arrangement to complement the description here.

We write the equation of the best-fit plane as $\mathbf{r} \cdot \hat{\mathbf{n}} = d$, with $\mathbf{r} = x\mathbf{i} + y\mathbf{j} + z\mathbf{k}$ being the vector position of a point in the plane, using the usual notation for the unit vectors in the Cartesian component directions, $\hat{\mathbf{n}} = l\mathbf{i} + m\mathbf{j} + n\mathbf{k}$ being a unit normal to the plane, and d being the distance of the plane from the origin. The equation of the plane then becomes $lx + my + nz = d$, which can be written as $z = (d - lx - my)/n$, or $z = \alpha - \beta x - \gamma y$, with $\alpha = d/n$, $\beta = l/n$, $\gamma = m/n$. A least-squares fit of this model to the mean annual cycle points \bar{y}_j allows us to determine values for α , β and γ . A little analysis shows that this corresponds to solving the equations

$$\begin{pmatrix} 12 & -S_x & -S_y \\ S_x & -S_{xx} & -S_{xy} \\ S_y & -S_{xy} & -S_{yy} \end{pmatrix} \begin{pmatrix} \alpha \\ \beta \\ \gamma \end{pmatrix} = \begin{pmatrix} S_z \\ S_{xz} \\ S_{yz} \end{pmatrix}$$

for α , β and γ , where the S_\bullet values are sums of components of the \bar{y}_j , i.e. S_x , S_y , S_z are the sums of the x , y and z components, S_{xx} , S_{yy} are the sums of the squared x and y components and S_{xy} , S_{xz} and S_{yz} are the sums of the appropriate component products.

Given the values α , β and γ , we can calculate

$$n = (1 + \beta^2 + \gamma^2)^{-1/2},$$

and $l = \beta n$, $m = \gamma n$, $d = \alpha n$, and can then construct the unit normal to the best-fit plane, $\hat{\mathbf{n}} = l\mathbf{i} + m\mathbf{j} + n\mathbf{k}$. We now wish to find a rotation taking $\hat{\mathbf{n}}$ into \mathbf{k} (the unit vector in the z -direction), thus rotating the best-fit plane into the x - y plane.

The required rotation may be determined using Rodrigues’ rotation formula, which states that the result of rotating a vector \mathbf{v} through an angle θ about the axis defined by another vector \mathbf{u} is

$$\mathbf{v}' = \mathbf{v} \cos \theta + \mathbf{u} \times \mathbf{v} \sin \theta + \mathbf{u}(\mathbf{u} \cdot \mathbf{v})(1 - \cos \theta).$$

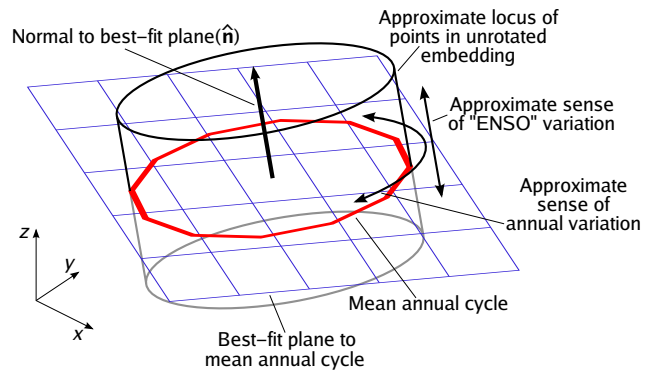


Fig. A1. Geometry of 3-D Isomap component rotation. The overall view is of an unrotated 3-D Isomap embedding. The thick red curve shows the mean annual cycle, the blue grid shows the best-fit plane to the mean annual cycle and the black arrow the normal to this plane, which we seek to rotate into the direction of the z -coordinate axis. Also illustrated is the approximate locus of points in the unrotated Isomap embedding (cf. Fig. 8a, for instance) and the approximate directions of annual (about the cylindrical locus) and ENSO variability (along the axis of the cylinder, orthogonal to the annual variation).

In the case here, we define a suitable rotation axis as $\mathbf{u} = \hat{\mathbf{n}} \times \mathbf{k} / |\hat{\mathbf{n}} \times \mathbf{k}|$, and the angle of rotation is $\theta = \cos^{-1}(\hat{\mathbf{n}} \cdot \mathbf{k})$ – this rotation will take $\hat{\mathbf{n}}$ into \mathbf{k} by rotating about a direction orthogonal to both $\hat{\mathbf{n}}$ and \mathbf{k} . Some simple algebra yields expressions for the individual rotated components:

$$\begin{aligned} x' &= nx - lz + \frac{mx - ly}{1 + n}m, \\ y' &= ny - mz - \frac{mx - ly}{1 + n}l, \\ z' &= nz + my + lx. \end{aligned}$$

Note that the rotation determined by Rodrigues’ formula is not unique. There remains an arbitrary phase to the annual cycle associated with rotations about the z -axis. For our purposes, this non-uniqueness is of no consequence – all we require is some rotation that will, as far as possible, unmix variations associated with the annual cycle from Isomap component #3 to reveal the interannual variability.

To see that the procedure we describe does indeed achieve this goal, see Fig. A2, where we display power spectra for Isomap components #1–#3 for k -Isomap results ($k=7$) using the ERSST v2 observational SST dataset. Figure A2a shows spectra for the raw Isomap components as calculated using Eq. (1) of Sect. 4.1.2. Here, there is a strong component at the annual frequency in all three components. Figure A2b shows spectra for the rotated Isomap components. The suppression of the annual signal in the spectrum of rotated Isomap component #3 is clear. Along with the high correlation between the rotated Isomap component #3 and the NINO3 SST index, this indicates that the Isomap algorithm successfully separates the annual cycle and ENSO variability out of the original SST field.

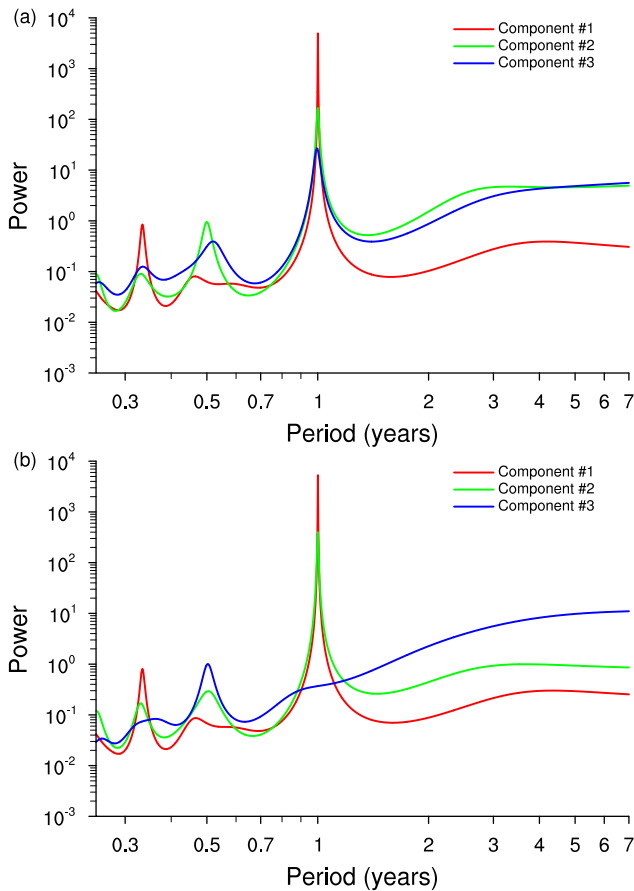


Fig. A2. Power spectra for Isomap components #1–3 from a raw SST k -Isomap analysis of the ERSST v2 observational SST dataset, showing the original Isomap output (a) and the rotated components (b).

A2 Four-dimensional case

The situation for four-dimensional embeddings is significantly more complicated than the three-dimensional case. This is due both to the more complex structure of the four-dimensional rotation group, $SO(4)$, compared to $SO(3)$, and to the absence of any easy geometrical intuition in four dimensions.

However, for the purposes of unmixing annual cycle variations from components #3 and #4 of a four-dimensional Isomap embedding, there are two observations that simplify matters considerably. (In the following, we denote the unit vectors in the coordinate directions for four-dimensional Euclidean space by $(\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3, \mathbf{e}_4)$.) The first observation is that any three-dimensional rotation is also a valid four-dimensional rotation, i.e. there are proper subgroups of $SO(4)$ that are isomorphic to $SO(3)$. If we have a matrix \mathbf{M}

representing an element of $SO(3)$, i.e.

$$\mathbf{M} = \begin{pmatrix} m_{11} & m_{12} & m_{13} \\ m_{21} & m_{22} & m_{23} \\ m_{31} & m_{32} & m_{33} \end{pmatrix}$$

with $\mathbf{M}^T \mathbf{M} = \mathbf{I}$ and $\det \mathbf{M} = 1$, then we can construct inclusion maps from $SO(3)$ into $SO(4)$ as

$$\mathbf{M}_3 = \begin{pmatrix} m_{11} & m_{12} & m_{13} & 0 \\ m_{21} & m_{22} & m_{23} & 0 \\ m_{31} & m_{32} & m_{33} & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \text{ and } \mathbf{M}_4 = \begin{pmatrix} m_{11} & m_{12} & 0 & m_{13} \\ m_{21} & m_{22} & 0 & m_{23} \\ 0 & 0 & 1 & 0 \\ m_{31} & m_{32} & 0 & m_{33} \end{pmatrix}.$$

These four-dimensional rotations represent rotations in the three-dimensional spaces spanned by $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_3\}$ and $\{\mathbf{e}_1, \mathbf{e}_2, \mathbf{e}_4\}$ respectively.

Secondly, since rotations by \mathbf{M}_3 do not affect the \mathbf{e}_4 component of any points and rotations by \mathbf{M}_4 do not affect the \mathbf{e}_3 components, we can compose rotations of these two types to unmix the annual variability from Isomap components #3 and #4 independently. Our approach is thus to use the three-dimensional rotation procedure described in Sect. A1 for each of Isomap components #3 and #4 in turn, so as to unmix annual variability from both of these components.

There is a caveat that should be applied to this procedure. As in the three-dimensional case, the rotations we use to unmix the annual variability from Isomap components #3 and #4 are not unique, and there is still a phase ambiguity present in both of the rotated components. Specifically, rotations leaving the \mathbf{e}_1 – \mathbf{e}_2 plane invariant will not affect the unmixing of the annual variability from the rotated Isomap components. Such rotations, represented by rotation matrices of the form

$$\mathbf{M}' = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & \cos \phi & -\sin \phi \\ 0 & 0 & \sin \phi & \cos \phi \end{pmatrix}$$

where ϕ is the rotation angle, do not alter the relationship between the rotated components #3 and #4 and the annual cycle components (#1 and #2), but they do alter the relative phasing between the rotated components #3 and #4. In practice, what this means is that, if one wishes to identify rotated Isomap components #3 and #4 as the “NINO3” and “WWV” components of ENSO variability, there is no guarantee that either of the rotated components is purely one form of ENSO variability or the other. This makes interpretation of the correlation results rather difficult. We have explored a number of approaches to unmixing the variability of these different degrees of freedom in this context, but there does not appear to be an easy a priori way to determine the angle ϕ to completely unmix the components. One possibility would be to rotate so as to maximise the correlations between rotated Isomap component #3 and the NINO3 SST index and

between rotated Isomap component #4 and the WWV time series, but this seems to be a rather unsatisfactorily ad hoc approach. These difficulties clearly have some bearing on interpretation of the results on correlations between the rotated component #4 and WWV reported in Sect. 5.1.

Acknowledgements. We would like to extend our thanks to three anonymous reviewers who provided extremely detailed and helpful comments. We acknowledge the modelling groups, the Program for Climate Model Diagnosis and Intercomparison (PCMDI) and the WCRP's Working Group on Coupled Modelling (WGCM) for their roles in making available the WCRP CMIP3 multi-model dataset. Support of this dataset is provided by the Office of Science, U.S. Department of Energy. NOAA ERSST v2 data were provided by the NOAA/OAR/ESRL PSD, Boulder, Colorado, USA, from their Web site at <http://www.cdc.noaa.gov/>.

Edited by: H. A. Dijkstra

Reviewed by: three anonymous referees

References

- AchutaRao, K. and Sperber, K. R.: Simulation of the El Niño Southern Oscillation: Results from the Coupled Model Intercomparison Project, *Clim. Dynam.*, 19, 191–209, 2002.
- AchutaRao, K. and Sperber, K. R.: ENSO simulation in coupled ocean-atmosphere models: are the current models better?, *Clim. Dynam.*, 27, 1–15, 2006.
- Aho, A. V., Hopcroft, J. E., and Ullman, J. D.: *Data Structures and Algorithms*, Addison-Wesley, 1983.
- An, S. I.: Interdecadal changes in the El Niño–La Niña asymmetry, *Geophys. Res. Lett.*, 31, 2617–2627, 2004.
- An, S. I. and Jin, F. F.: Nonlinearity and asymmetry of ENSO, *J. Climate*, 17, 2399–2412, 2004.
- An, S. I., Ham, Y. G., Kug, J. S., Jin, F. F., and Kang, I. S.: El Niño–La Niña asymmetry in the Coupled Model Intercomparison Project simulations, *J. Climate*, 18, 2617–2627, 2005a.
- An, S. I., Hsieh, W. W., and Jin, F. F.: A nonlinear analysis of the ENSO cycle and its interdecadal changes, *J. Climate*, 18, 3229–3239, 2005b.
- Belkin, M. and Niyogi, P.: Laplacian eigenmaps for dimensionality reduction and data representation, *Neural Comput.*, 15, 1373–1396, 2003.
- Bishop, C. M., Svensen, M., and Williams, C. K. I.: GTM: The generative topographic mapping, *Neural Comput.*, 10, 215–234, 1998.
- Bollt, E.: Attractor modeling and empirical nonlinear model reduction of dissipative dynamical systems, *Int. J. Bifurc. Chaos*, 17, 1199–1219, 2007.
- Borg, I. and Groenen, P.: *Modern Multidimensional Scaling: Theory and Applications*, Springer Series in Statistics, Springer-Verlag, London, 1997.
- Broomhead, D. and Kirby, M.: Dimensionality reduction using secant-based projection methods: The induced dynamics in projected systems, *Nonlin. Dyn.*, 41, 47–67, 2005.
- Burgers, G.: The El Niño stochastic oscillator, *Clim. Dynam.*, 15, 521–531, 1999.
- Collins, W. D., Bitz, C. M., Blackmon, M. L., Bonan, G. B., Bretherton, C. S., Carton, J. A., Chang, P., Doney, S. C., Hack, J. J., Henderson, T. B., Kiehl, J. T., Large, W. G., McKenna, D. S., Santer, B. D., and Smith, R. D.: The Community Climate System Model version 3 (CCSM3), *J. Climate*, 19, 2122–2143, 2006.
- de Silva, V. and Tenenbaum, J. B.: Global versus local methods in nonlinear dimensionality reduction, in: *Advances in Neural Information Processing Systems*, edited by: Becker, S., Thrun, S., and Obermayer, K., 15, 705–712, MIT Press, Cambridge, 2002.
- Delworth, T. L., Broccoli, A. J., Rosati, A., Stouffer, R. J., Balaji, V., Beesley, J. A., Cooke, W. F., Dixon, K. W., Dunne, J., Dunne, K. A., Durachta, J. W., Findell, K. L., Ginoux, P., Gnanadesikan, A., Gordon, C. T., Griffies, S. M., Gudgel, R., Harrison, M. J., Held, I. M., Hemler, R. S., Horowitz, L. W., Klein, S. A., Knutson, T. R., Kushner, P. J., Langenhorst, A. R., Lee, H. C., Lin, S. J., Lu, J., Malyshev, S. L., Milly, P. C. D., Ramaswamy, V., Russell, J., Schwarzkopf, M. D., Shevliakova, E., Sirutis, J. J., Spelman, M. J., Stern, W. F., Winton, M., Wittenberg, A. T., Wyman, B., Zeng, F., and Zhang, R.: GFDL's CM2 global coupled climate models. Part I: Formulation and simulation characteristics, *J. Climate*, 19, 643–674, 2006.
- Donoho, D. L. and Grimes, C.: Hessian eigenmaps: Locally linear embedding techniques for high-dimensional data, *Proc. Natl. Acad. Sci. USA*, 100, 5591–5596, 2003.
- Donoho, D. L. and Grimes, C.: Image manifolds which are isometric to Euclidean space, *J. Math. Imaging Vis.*, 23, 5–24, 2005.
- Fedorov, A. V. and Philander, S. G.: Is El Niño changing?, *Science*, 288, 1997–2002, 2000.
- Furevik, T., Bentsen, M., Drange, H., Kindem, I. K. T., Kvamstø, N. G., and Sorteberg, A.: Description and evaluation of the Bergen climate model: ARPEGE coupled with MICOM, *Clim. Dynam.*, 21, 27–51, 2003.
- Gámez, A. J., Zhou, C. S., Timmermann, A., and Kurths, J.: Nonlinear dimensionality reduction in climate data, *Nonlin. Processes Geophys.*, 11, 393–398, 2004, <http://www.nonlin-processes-geophys.net/11/393/2004/>.
- Gordon, C., Cooper, C., Senior, C. A., Banks, H., Gregory, J. M., Johns, T. C., Mitchell, J. F. B., and Wood, R. A.: The simulation of SST, sea ice extents and ocean heat transports in a version of the Hadley Centre coupled model without flux adjustments, *Clim. Dynam.*, 16, 147–168, 2000.
- Gordon, H. B., Rotstayn, L. D., McGregor, J. L., Dix, M. R., Kowalczyk, E. A., O'Farrell, S. P., Waterman, L. J., Hirst, A. C., Wilson, S. G., Collier, M. A., Watterson, I. G., and Elliott, T. I.: The CSIRO Mk3 Climate System Model, Tech. Rep. 60, CSIRO Atmospheric Research, Aspendale, 2002.
- Gower, J. C.: Some distance properties of latent root and vector methods used in multivariate analysis, *Biometrika*, 53, 325–338, 1966.
- Guilyardi, E.: El Niño-mean state-seasonal cycle interactions in a multi-model ensemble, *Clim. Dynam.*, 26, 329–348, 2006.
- Hannachi, A., Jolliffe, I. T., and Stephenson, D. B.: Empirical orthogonal functions and related techniques in atmospheric science: A review, *Int. J. Climatol.*, 27, 1119–1152, 2007.
- Heroux, M. A., Bartlett, R. A., Howle, V. E., Hoekstra, R. J., Hu, J. J., Kolda, T. G., Lehoucq, R. B., Long, K. R., Pawlowski, R. P., Phipps, E. T., Salinger, A. G., Thornquist, H. K., Tuminaro, R. S., Willenbring, J. M., Williams, A., and Stanley, K. S.: An overview of the Trilinos Project, *ACM T. Math. Software*, 31, 397–423, 2005.

- Hinton, G. E. and Salakhutdinov, R. R.: Reducing the dimensionality of data with neural networks, *Science*, 313, 504–507, 2006.
- Hsieh, W. W.: Nonlinear multivariate and time series analysis by neural network methods, *Rev. Geophys.*, 42, 2617–2627, 2004.
- Jin, F. F., Neelin, J. D., and Ghil, M.: El Niño on the Devil's Staircase – Annual subharmonic steps to chaos, *Science*, 264, 70–72, 1994.
- Johns, T. C., Durman, C. F., Banks, H. T., Roberts, M. J., McLaren, A. J., Ridley, J. K., Senior, C. A., Williams, K. D., Jones, A., Rickard, G. J., Cusack, S., Ingram, W. J., Crucifix, M., Sexton, D. M. H., Joshi, M. M., Dong, B. W., Spencer, H., Hill, R. S. R., Gregory, J. M., Keen, A. B., Pardaens, A. K., Lowe, J. A., Bodas-Salcedo, A., Stark, S., and Searl, Y.: The new Hadley Centre Climate Model (HadGEM1): Evaluation of coupled simulations, *J. Climate*, 19, 1327–1353, 2006.
- K-1 model developers: K-1 coupled GCM (MIROC) description, Tech. Rep. 1, Centre for Climate System Research, University of Tokyo, 2004.
- Kessler, W. S.: Is ENSO a cycle or a series of events?, *Geophys. Res. Lett.*, 29, 719–730, 2002.
- Kim, S. J., Flato, G. M., Boer, G. J., and McFarlane, N. A.: A coupled climate model simulation of the Last Glacial Maximum, Part 1: transient multi-decadal response, *Clim. Dynam.*, 19, 515–537, 2002.
- Latif, M., Sperber, K., Arblaster, J., Braconnot, P., Chen, D., Colman, A., Cubasch, U., Cooper, C., Delecluse, P., DeWitt, D., Fairhead, L., Flato, G., Hogan, T., Ji, M., Kimoto, M., Kitoh, A., Knutson, T., Le Treut, H., Li, T., Manabe, S., Marti, O., Mechoso, C., Meehl, G., Power, S., Roeckner, E., Sirven, J., Terray, L., Vintzileos, A., Voss, R., Wang, B., Washington, W., Yoshikawa, I., Yu, J., and Zebiak, S.: ENSIP: the El Niño simulation intercomparison project, *Clim. Dynam.*, 18, 255–276, 2001.
- Leloup, J. A., Lachkar, Z., Boulanger, J. P., and Thiria, S.: Detecting decadal changes in ENSO using neural networks, *Clim. Dynam.*, 28, 147–162, 2007.
- Lin, T., Zha, H. B., and Lee, S. U.: Riemannian manifold learning for nonlinear dimensionality reduction, *Lect. Notes Comput. Sci.*, 3951, 44–55, 2006.
- Marti, O., Braconnot, P., Bellier, J., Benschila, R., Bony, S., Brockmann, P., Cadulle, P., Caubel, A., Denvil, S., Dufresne, J. L., Fairhead, L., Filiberti, M.-A., Fichet, T., Friedlingstein, P., Grandpeix, J.-Y., Hourdin, F., Krinner, G., Lévy, C., Musat, I., and Talandier, C.: The new IPSL climate system model: IPSL-CM4, Tech. rep., Institut Pierre Simon Laplace des Sciences de l'Environnement Global, IPSL, Case 101, 4 place Jussieu, Paris, France, 2005.
- McPhaden, M. J.: Climate oscillations – Genesis and evolution of the 1997-98 El Niño, *Science*, 283, 950–954, 1999.
- McPhaden, M. J.: Tropical Pacific Ocean heat content variations and ENSO persistence barriers, *Geophys. Res. Lett.*, 30, 719–730, 2003.
- McPhaden, M. J., Zebiak, S. E., and Glantz, M. H.: ENSO as an integrating concept in Earth science, *Science*, 314, 1740–1745, 2006.
- Mechoso, C. R., Robertson, A. W., Barth, N., Davey, M. K., Delecluse, P., Gent, P. R., Ineson, S., Kirtman, B., Latif, M., Le Treut, H., Nagai, T., Neelin, J. D., Philander, S. G. H., Polcher, J., Schopf, P. S., Stockdale, T., Suarez, M. J., Terray, L., Thual, O., and Tribbia, J. J.: The seasonal cycle over the tropical Pacific in coupled ocean-atmosphere general-circulation models, *Mon. Weather Rev.*, 123, 2825–2838, 1995.
- Meinen, C. S. and McPhaden, M. J.: Observations of warm water volume changes in the equatorial Pacific and their relationship to El Niño and La Niña, *J. Climate*, 13, 3551–3559, 2000.
- Min, S. K., Legutke, S., Hense, A., and Kwon, W. T.: Internal variability in a 1000-yr control simulation with the coupled climate model ECHO-G – I. Near-surface temperature, precipitation and mean sea level pressure, *Tellus A*, 57, 605–621, 2005.
- Monahan, A. H.: Nonlinear principal component analysis: Tropical Indo-Pacific sea surface temperature and sea level pressure, *J. Climate*, 14, 219–233, 2001.
- Monahan, A. H. and Dai, A. G.: Spatial and temporal structure of ENSO nonlinearity, *J. Climate*, 17, 3026–3036, 2004.
- Moore, A. M. and Kleeman, R.: Stochastic forcing of ENSO by the intraseasonal oscillation, *J. Climate*, 12, 1199–1220, 1999.
- Neelin, J. D., Latif, M., Allaart, M. A. F., Cane, M. A., Cubasch, U., Gates, W. L., Gent, P. R., Ghil, M., Gordon, C., Lau, N. C., Mechoso, C. R., Meehl, G. A., Oberhuber, J. M., Philander, S. G. H., Schopf, P. S., Sperber, K. R., Sterl, A., Tokioka, T., Tribbia, J., and Zebiak, S. E.: Tropical air-sea interaction in general-circulation models, *Clim. Dynam.*, 7, 73–104, 1992.
- North, G. R.: Empirical orthogonal functions and normal modes, *J. Atmos. Sci.*, 41, 879–887, 1984.
- Philander, S. G.: El Niño, La Niña and the Southern Oscillation, Academic Press, San Diego, 1990.
- Press, W. H., Teukolsky, S. A., Vetterling, W. T., and Flannery, B. P.: Numerical Recipes in C: The Art of Scientific Computing, Cambridge University Press, 2nd. edn., 1992.
- Roweis, S. T. and Saul, L. K.: Nonlinear dimensionality reduction by locally linear embedding, *Science*, 290, 2323–2326, 2000.
- Salas-Méla, D., Chauvin, F., Déqué, M., Douville, H., Guérémy, J. F., Marquet, P., Planton, S., Royer, J. F., and Tyteca, S.: Description and validation of the CNRM-CM3 global coupled model, CNRM working note 103, 2005.
- Schmidt, G. A., Ruedy, R., Hansen, J. E., Aleinov, I., Bell, N., Bauer, M., Bauer, S., Cairns, B., Canuto, V., Cheng, Y., Genio, A. D., Faluvegi, G., Friend, A. D., Hall, T. M., Hu, Y. Y., Kelley, M., Kiang, N. Y., Koch, D., Lacis, A. A., Lerner, J., Lo, K. K., Miller, R. L., Nazarenko, L., Oinas, V., Perlwitz, J., Perlwitz, J., Rind, D., Romanou, A., Russell, G. L., Sato, M., Shindell, D. T., Stone, P. H., Sun, S., Tausnev, N., Thresher, D., and Yao, M. S.: Present-day atmospheric simulations using GISS ModelE: Comparison to in situ, satellite, and reanalysis data, *J. Climate*, 19, 153–192, 2006.
- Smith, T. M. and Reynolds, R. W.: Improved extended reconstruction of SST (1854–1997), *J. Climate*, 17, 2466–2477, 2004.
- Tenenbaum, J. B., de Silva, V., and Langford, J. C.: A global geometric framework for nonlinear dimensionality reduction, *Science*, 290, 2319–2323, 2000.
- Thompson, C. J. and Battisti, D. S.: A linear stochastic dynamical model of ENSO. Part I: Model development, *J. Climate*, 13, 2818–2832, 2000.
- Torgerson, W. S.: Multidimensional scaling: I. Theory and method, *Psychometrika*, 17, 401–419, 1952.
- Tziperman, E., Stone, L., Cane, M. A., and Jarosh, H.: El Niño chaos: Overlapping of resonances between the seasonal cycle and the Pacific ocean-atmosphere oscillator, *Science*, 264, 72–74, 1994.

- UK Meteorological Office: HadISST 1.1 – Global sea-Ice coverage and SST (1870-Present), Internet, available from <http://badc.nerc.ac.uk/data/hadisst/>, 2006.
- van Oldenborgh, G. J., Philip, S. Y., and Collins, M.: El Niño in a changing climate: a multi-model study, *Ocean Sci.*, 1, 81–95, 2005, <http://www.ocean-sci.net/1/81/2005/>.
- Volodin, E. M. and Diansky, A. N.: El Niño reproduction in a coupled general circulation model of atmosphere and ocean, *Russ. Meteorol. Hydrol.*, 12, 5–14, 2004.
- von Storch, H. and Zwiers, F. W.: *Statistical Analysis in Climate Research*, Cambridge University Press, 2003.
- Weinberger, K. Q. and Saul, L. K.: Unsupervised learning of image manifolds by semidefinite programming, *Int. J. Comput. Vision*, 70, 77–90, 2006.
- Wu, A. and Hsieh, W. W.: Nonlinear interdecadal changes of the El Niño-Southern Oscillation, *Clim. Dynam.*, 21, 719–730, 2003.
- Yu, Y. Q., Zhang, X. H., and Guo, Y. F.: Global coupled ocean-atmosphere general circulation models in LASG/IAP, *Adv. Atmos. Sci.*, 21, 444–455, 2004.
- Yukimoto, S., Noda, A., Kitoh, A., Hosaka, M., Yoshimura, H., Uchiyama, T., Shibata, K., Arakawa, O., and Kusunoki, S.: Present-day climate and climate sensitivity in the Meteorological Research Institute coupled GCM version 2.3 (MRI-CGCM2.3), *J. Meteorol. Soc. Jpn.*, 84, 333–363, 2006.
- Zebiak, S. E. and Cane, M. A.: A model El Niño Southern Oscillation, *Mon. Weather Rev.*, 115, 2262–2278, 1987.