Max-Planck-Institut
für Meteorologie

REPORT No. 112

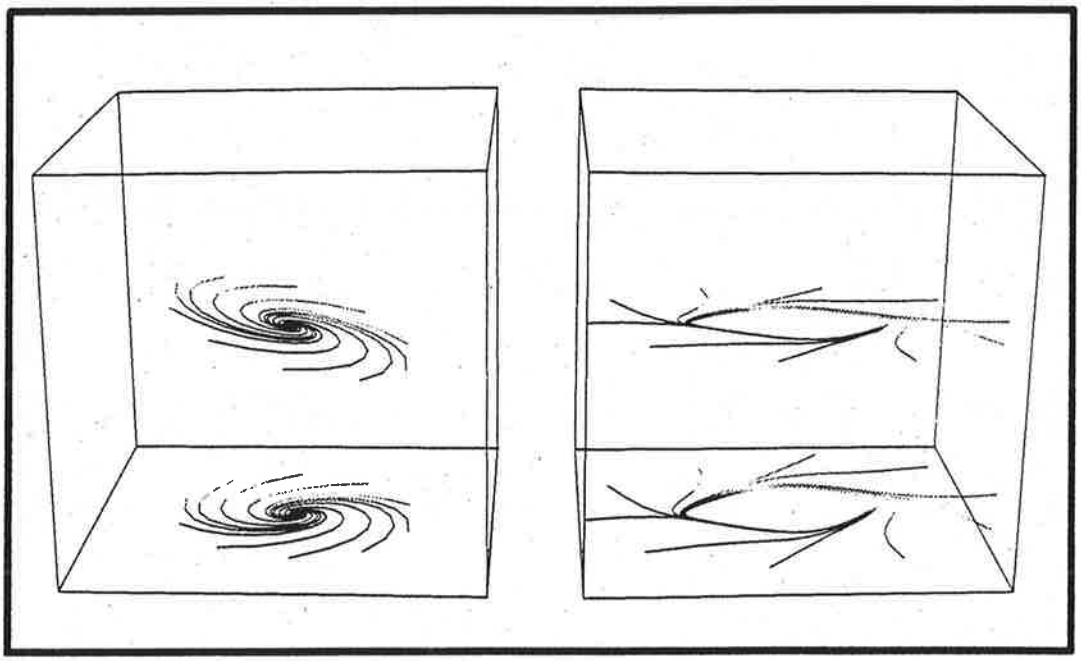

RECONSTRUCTION OF THE EL NIÑO ATTRACTOR WITH NEURAL NETWORKS

von

BJÖRN GRIEGER • MOJIB LATIF

HAMBURG, SEPTEMBER 1993

AUTHORS:

Björn Grieger

Max–Planck–Institut
für Meteorologie

present affilliation:
Universität Bremen
FB 5 – Geowissenschaften
Postfach 33 04 40
28334 Bremen

Mojib Latif

Max–Planck–Institut
für Meteorologie

MAX–PLANCK–INSTITUT
FÜR METEOROLOGIE
BUNDESSTRASSE 55
D–20146 Hamburg
F.R. GERMANY

Tel.: +49–(0)40–4 11 73–0
Telex: 211092 mpime d
Telemail: MPI.METEOROLOGY
Telefax: +49–(0)40–4 11 73–298

REPb 112

ISSN 0937–1060

Reconstruction of the El Niño Attractor with Neural Networks

B. Grieger M. Latif

July 20, 1993

Abstract

Based on a combined data set of sea surface temperature, zonal surface wind stress and upper ocean heat content the dynamics of the El Niño phenomenon is investigated. In a reduced phase space spanned by the first four EOFs two different stochastic models are estimated from the data. A nonlinear model represented by a simulated neural network is compared with a linear model obtained with the Principal Oscillation Pattern (POP) analysis. While the linear model is limited to damped oscillations onto a fix point attractor, the nonlinear model recovers a limit cycle attractor. This indicates that the real system is located above the bifurcation point in parameter space supporting self-sustained oscillations. The results are discussed with respect to consistency with current theory.

1 Introduction

A major task in climate research and modeling is to reduce the number of degrees of freedom from observational or model produced data sets. Several statistical methods have been developed for this purpose, such as real or complex Empirical Orthogonal Functions (EOFs) or Canonical Correlation Analysis (CCA). More recently Hasselmann (1988) proposed the method of Principal Interaction Patterns (PIPs) to derive simplified dynamical models from large data sets. In the linear case the PIPs reduce to the Principal Oscillation Patterns (POPs) (Storch *et al.* (1988,1990)) which represent the damped oscillations of a stochastically driven system.

During the last few years neural networks have successfully been applied to a variety of problems. It has been shown by Elsner and Tsonis (1992) that neural networks can be used successfully for the prediction of a univariate chaotic time series. Tang (1992) trained neural networks to forecast the El Niño state for a certain lead time and demonstrated that its skill is comparable to or better than that of a linear statistical model.

Here we use neural networks to get more insight into the dynamics of the El Niño/Southern Oscillation (ENSO) phenomenon (Rasmusson and Carpenter (1982), Cane *et al.* (1986), Philander (1990)). ENSO is the most prominent climate signal on the short-range climatic time scale and recent observational and modeling studies suggest that ENSO is based on a low-frequency oscillation (Schopf and Suarez (1988), Graham and White (1988), Cane *et al.* (1990), Chao and Philander (1991), Latif *et al.* (1993)). A conceptual model has been developed, the so-called "delayed action oscillator" (Schopf and Suarez (1988)) which is commonly regarded as a paradigm for ENSO and is based on equatorial

wave propagation and reflection (Zebiak and Cane (1987), Battisti (1988), Battisti and Hirst (1989)). Neelin and Jin (1993) showed that the "delayed action oscillator" can be regarded as an extreme case of a general mixed surface/subsurface dynamics mode. We are here concerned with the question of what the underlying dynamics of the ENSO cycle is and to what extent this dynamics is consistent with current theory.

Our primary tool is the neural network technique (see e. g. Widrow and Lehr (1990)). The optimization of a suitable structured neural network with the backpropagation learning algorithm (Rumelhart *et al.* (1986)) is comparable with fitting a parameterized function by error minimization using steepest descent methods. The neural network formulation provides a large class of model functions accompanied by analytical derivatives, which strongly accelerates the optimization procedure.

This paper is organized as follows. Section 2 deals with the description of the data whose degrees of freedom have been reduced by applying EOF analysis. In section 3 we give a brief description of the neural network approach. In section 4 we derive reduced state space models of ENSO based on a linear (POP) and the (non-linear) neural network approach. The relationship of our results to the ENSO theory is discussed in section 5. We conclude the paper in section 6.

2 Data and EOF analysis

We use bimonthly observations of sea surface temperature, depth of the 20°C-isotherm (a measure of upper ocean heat content), and zonal surface wind stress for the period 1967 to 1986. The same data set was used by Latif *et al.* (1993), who performed a POP

analysis of this combined data set. We removed prior to the analysis the annual cycle and linear trend from the data. Each quantity was normalized by its spatially averaged standard deviation so that they all have the same weight.

We then performed an EOF analysis of the data. The first four EOFs, explaining 23.3%, 9.6%, 4.6%, and 4.0% of the variance in the data, respectively, are shown in Fig. 1. The first two combined EOFs are associated with the ENSO phenomenon. As shown by Latif *et al.* (1993) and below, the dominant POP mode of the data is mainly composed of the two leading EOFs. The first combined EOF represents conditions during the extreme phase of ENSO, with anomalously warm surface waters in the eastern equatorial Pacific, a drop in upper ocean heat content in the western Pacific, and westerly wind stress anomalies centered near the dateline. The second combined EOF is dominated by a strong heat content anomaly centered at the equator. As discussed by Latif *et al.* (1993) these two patterns describe well the dominant mode of interannual variability in the tropical Pacific and are consistent with the conceptual model of "delayed action oscillation" (Schopf and Suarez (1988)). According to this picture, the propagation of equatorial waves and their reflection at meridional boundaries are crucial in maintaining the ENSO cycle. (For a more detailed discussion of the "delayed action oscillator" the reader is referred to Graham and White (1988), Chao and Philander (1991), Cane *et al.* (1990).)

The purpose of the EOF analysis is the reduction of the number of degrees of freedom. Various methods have been described to determine the truncation point (Richman *et al.*, 1992), i. e. the number of EOFs which are kept for further analysis. None of the proposed criteria take into account the dynamics, i. e. the temporal sequence of the patterns. Here

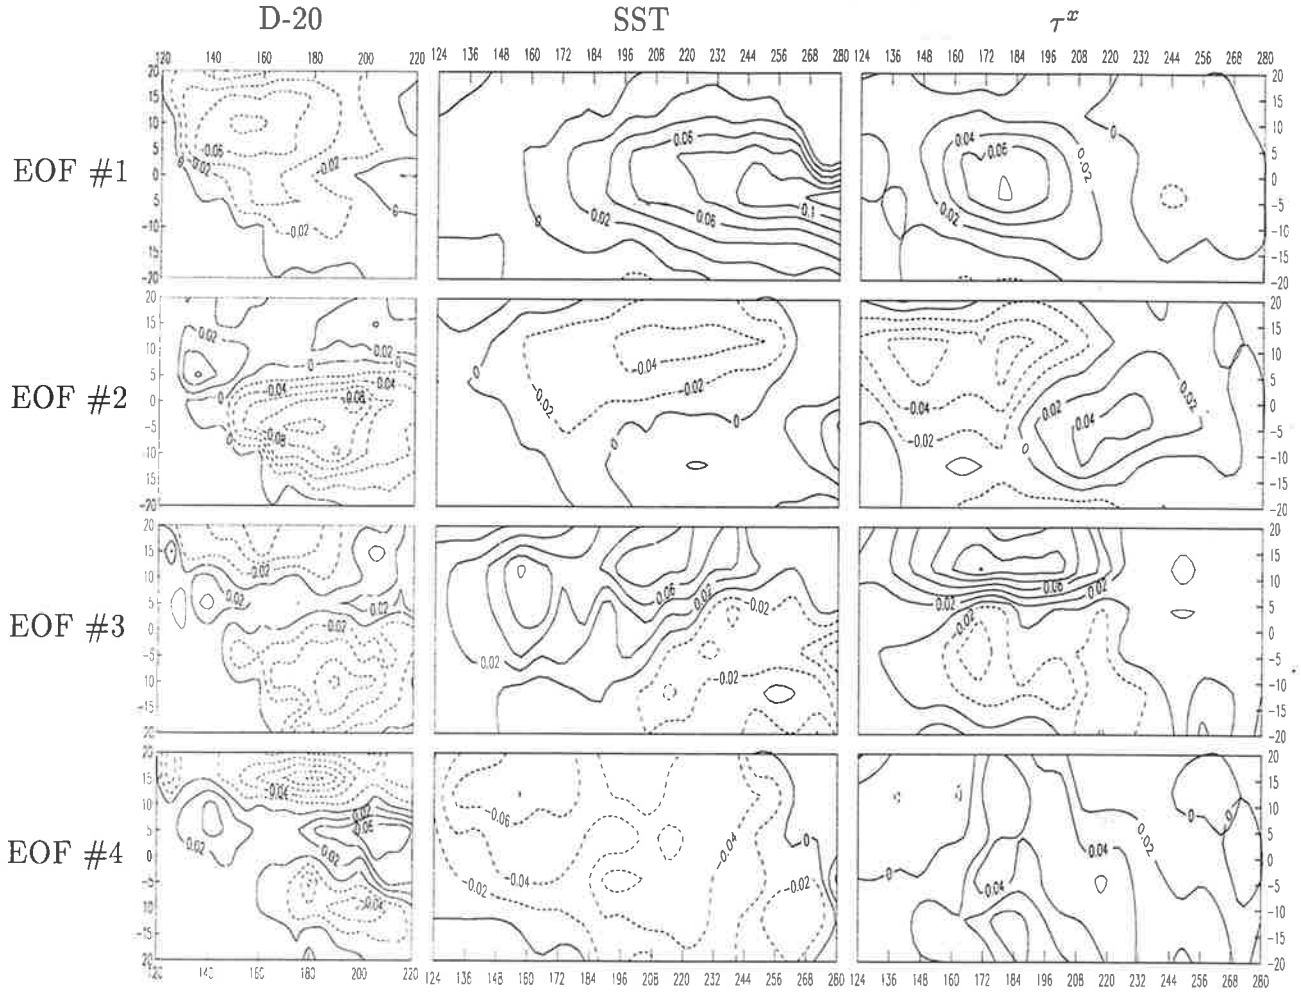

Figure 1: The first four EOFs (from top to bottom) of the data. *Left:* Heat content (depth of the 20°C-isotherm). *Center:* Sea surface temperature. *Right:* Zonal wind stress. Longitude is in degrees east, latitude is in degrees north (positive) and south (negative).

we use a criterion which measures the "consistency" of the trajectory in EOF space.

Given an evenly spaced time series of state space vectors $(\mathbf{x}(t))_{t=1,\ldots,T}$ we define the progression vectors

$$\Delta\mathbf{x}(t) := \mathbf{x}(t+1) - \mathbf{x}(t). \tag{1}$$

A badness function which measures the inconsistency of the observed state space trajectory can be expressed through

$$\varepsilon = \sqrt{\frac{2}{(T-1)(T-2)} \sum_{t_1=1}^{T-1} \sum_{t_2=t_1+1}^{T-1} \frac{\|\Delta\mathbf{x}(t_1) - \Delta\mathbf{x}(t_2)\|^2}{\|\mathbf{x}(t_1) - \mathbf{x}(t_2)\|^2}}, \tag{2}$$

where the sum is taken over all pairs of observed points (normalized with the total number of pairs). The contribution of every pair to the badness becomes large when the two observed states are close in the state space and the progression vectors differ considerably. The idea is that if the state space trajectory is approximately deterministic the difference in the chance Δx of two states of the system should be small if the difference in the states is small. Stochastic forcing increases the badness by producing different progression vectors even for exactly equal states. So we are looking for the state space dimension in which the system's trajectory can be explained with minimum stochastic forcing.

Of course the badness function (2) is biased. An exact measure of the strength of the stochastic component would be the mean difference of progression vectors in the limit of vanishing distance in the state space, i. e.

$$\varepsilon' = \sqrt{\left\langle \lim_{\mathbf{x} \to \mathbf{x}_0} \|\Delta\mathbf{x} - \Delta\mathbf{x}_0\|^2 \right\rangle}, \tag{3}$$

where $\langle \ldots \rangle$ denotes averaging over the state space (over all $\mathbf{x}_0$) and all realisations, and $\Delta\mathbf{x}$ and $\Delta\mathbf{x}_0$ are the progression vectors at $\mathbf{x}$ and $\mathbf{x}_0$, respectively. But the estimation of

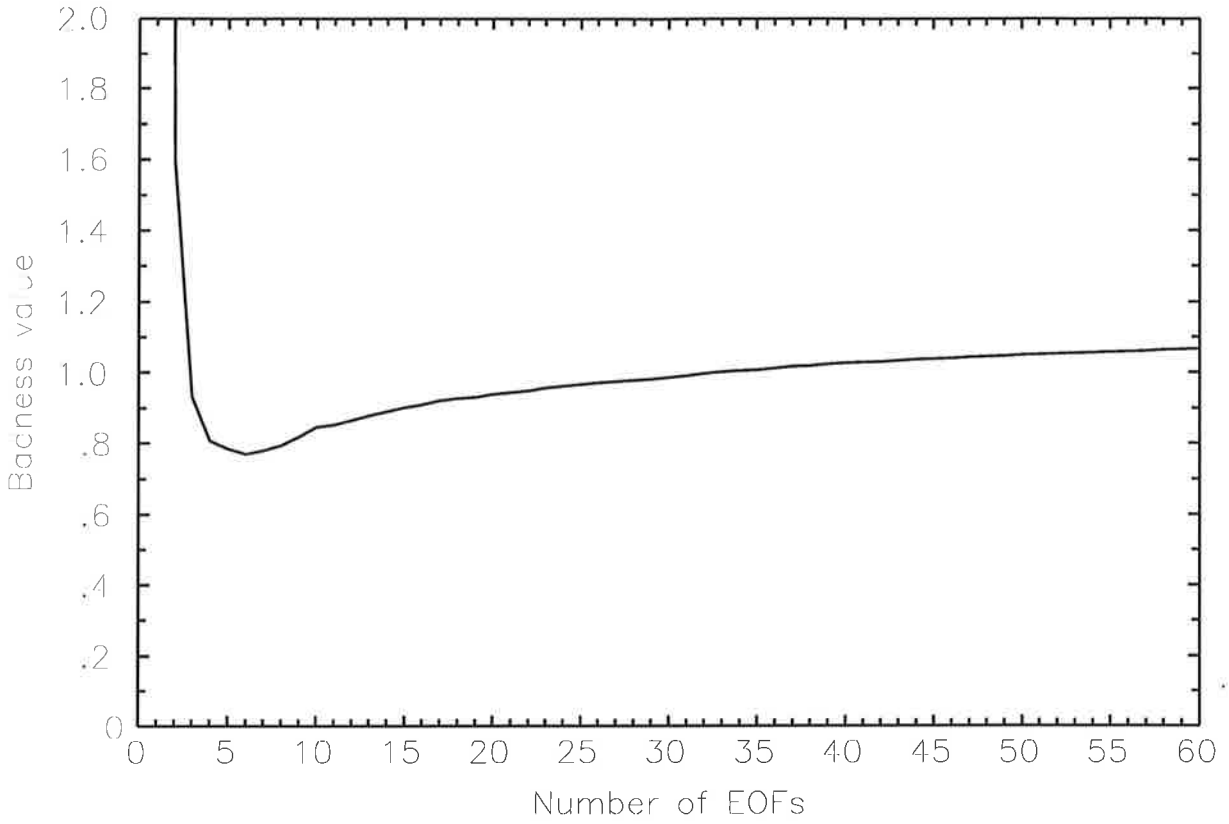

Figure 2: "Badness" (scatter of the observed phase space trajectory as a function of the number of EOFs).

this value requires a large amount of data which is not available in our case. Therefore we use Eq. (2) as a crude but reasonable proxy.

The value of the badness as a function of the number of used EOFs is shown in Fig. 2. The function has a minimum at six, but we decided to use only four EOFs, which explain 41.5% of the data variance. The increase in badness is not large, and with less degrees of freedom the model fitting procedure is accelerated. Moreover four dimensions can nicely be presented in a coloured three dimensional picture. In Fig. 3 the "raw" trajectory in the four-dimensional state space (spanned by the first four EOFs) is shown.

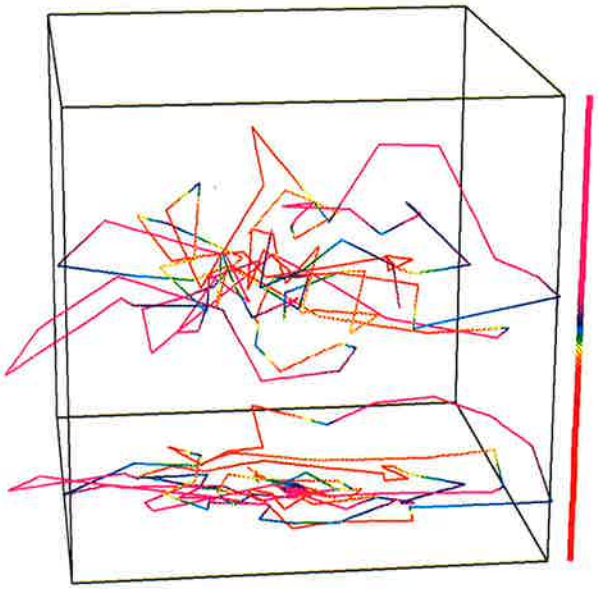

Figure 3: Observed trajectory in the four-dimensional state space. The spatial dimensions represent the first three EOFs, the color corresponds to the fourth EOF. The color scale with respect to the space scale is displayed at the right margin. (The curve at the bottom of the box is just a projection—a "shadow"—to elucidate the three dimensional structure.)

3 A neural network as non-linear function approximator

To estimate a non-linear system function we need a method of reconstructing the flow field at every point in the state space from observations of the flow field at a limited number of points, i. e. a model function depending on a few free parameters has to be fitted to the data. The neural network approach described here provides a suitable class of nonlinear functions accompanied by a matching optimization procedure.

In a layered feed-forward network (an example is shown in Fig. 4), the output of neuron i in layer k is given by

$$y_i^{(k)} = f_{\text{out}}^{(ki)} \left(\sum_{j=0}^{N_{k-1}} W_{ij}^{(k)} y_j^{(k-1)} \right), \tag{4}$$

where N_{k-1} is the number of neurons in layer $k-1$, $W_{ij}^{(k)}$ is the weight for input j of neuron i in layer k, and $f_{\text{out}}^{(ki)}$ is the output function of neuron i in layer k. In the present case a two-layered network with $f_{\text{out}}^{(1i)} = \tanh$ and $f_{\text{out}}^{(2i)} = \text{id}$ will be used. $y_j^{(0)}$ denotes the input vector. The so-called dummy inputs are $y_0^{(k)} = 1$. Given a training set of input vectors accompanied by desired output vectors, the neural network can be optimized by backpropagation of the output errors through the whole network to adjust all weights, see e. g. Rumelhart *et al.* (1986).

Consider a system represented by a state vector $\mathbf{x} = (x_1, \ldots, x_n)$ which varies with time. The temporal evolution of the system is described by

$$\mathbf{x}(t+1) = \mathbf{F}[\mathbf{x}(t); \alpha_1, \ldots, \alpha_m] + \text{noise}, \tag{5}$$

where $\mathbf{F}$ is a function which has to be prescribed but may depend on an arbitrary number

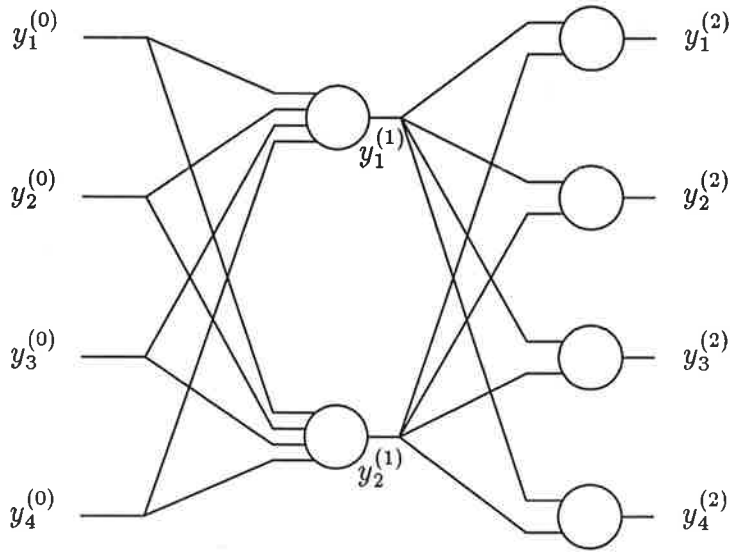

Figure 4: Structure of the neural network used for the estimation of the flow field in a four-dimensional state space. Input (left side) is an actual state space vector, output (right side) is the predicted state space vector for the next time step. The number of neurons in the hidden layer (two in the example shown) determines the complexity of the flow field.

of free parameters $(\alpha_j)_{j=1,\ldots,m}$. The function parameters $(\alpha_j)_{j=1,\ldots,m}$ can be obtained from a given set of observations $(\mathbf{x}(t))_{t=1,\ldots,T}$ by minimization of the lag 1 prediction error:

$$\sum_{t=1}^{T-1} \|\mathbf{x}(t+1) - \mathbf{F}[\mathbf{x}(t); \alpha_1, \ldots, \alpha_m]\|^2 = \min. \tag{6}$$

If we describe the problem in terms of neural networks, the system function $\mathbf{F}$ can be represented by a two-layered nonlinear network ($f_{\text{out}}^{(1i)} = \tanh$) where a state space vector $\mathbf{x}(t)$ at a certain time is the input vector and the state space vector $\mathbf{x}(t+1)$ one time step later is the corresponding output vector, i. e.

$$\mathbf{x}(t+1) = \mathbf{W}_0^{(2)} + \mathbf{W}^{(2)} \tanh[\mathbf{W}_0^{(1)} + \mathbf{W}^{(1)}\mathbf{x}(t)] + \text{noise}. \tag{7}$$

Here $\mathbf{W}^{(1)}$ and $\mathbf{W}^{(2)}$ denote the matrices of the weights $(W_{ij}^{(1)})_{i=1\ldots N_1, j=1\ldots n}$ and $(W_{ij}^{(2)})_{i=1\ldots n, j=1\ldots N_1}$ in the first and the second layer, respectively, and $\mathbf{W}_0^{(1)}$ and $\mathbf{W}_0^{(2)}$ denote the vectors of the weights $(W_{i0}^{(1)})_{i=1\ldots N_1}$ and $(W_{i0}^{(2)})_{i=1\ldots n}$ for the dummy inputs.

These weights are the function parameters $(\alpha_j)_{j=1,\ldots,m}$, and the total number of parameters is $m = n+2nN_1+N_1$, where n is the number of dimensions and N_1 is the number of hidden neurons. Fig. 4 shows the structure of a neural network for a four-dimensional system.

The weights of the network could be estimated by the following operations: An input vector $\mathbf{x}(t)$ is chosen (randomly or in cyclic order) from the training set and presented to the network. The desired output vector $\mathbf{x}(t+1)$ is used to adjust the weights of the network—the parameters $(\alpha_j)_{j=1,\ldots,m}$—with the backpropagation algorithm. This is repeated until convergence is achieved. Initially the network weights are set to small random values.

The number of parameters of the neural network model is $n + 2nN_1 + N_1$. With increasing number of neurons in the hidden layer, N_1, the number of parameters is increased and the lag 1 prediction error for the training data decreases more and more. But if the number of parameters is too large, the model tries to fit the stochastic forcing, and when applied to a new system state which was not incorporated in the training data, the prediction skill is low (compare Barnett and Hasselmann (1979)). To determine the optimal number of hidden neurons, only a part of the available data is used. The prediction error is then checked against the remaining data, as demonstrated below.

A general problem of the backpropagation training of neural networks is the large computation expense, which limits the size of the networks. We needed some hundred thousand presentations of the data before convergence was achieved. The number of necessary training cycles increases dramatically with the size of the network, so even a hardware implementation with completely parallel processing would not solve this problem. Some computation time could be saved if the steepest descent strategy which we use is replaced by a line-search algorithm, see e. g. Tang (1992).

We made some experiments with different starting values for the network weights which are ususally set to small random values. In a few cases the optimization procedure was trapped in a local minimum. Therefore we cannot be sure that the estimated models represent indeed the global minima of the target function. More sophisticated nonlinear optimization techniques such as simulated annealing are very time consuming and have not been applied to neural networks yet.

4 Reduced state space models

The raw state space trajectory of the data, which was shown in Fig. 3, is very noisy and illustrates the need for constructing reduced state space models.

First, we fitted a linear model—the POP model—to the data. POPs are the eigenvectors of the system matrix obtained by fitting a multi-variate first order autoregressive (Markov) process to the data. POPs are generally complex, their complex amplitudes satisfying the standard damped harmonic oscillator equation (Hasselmann (1988), Storch *et al.* (1988)). As described previously by Latif *et al.* (1993), the dominant POP mode is consistent with the "delayed action oscillator" scenario. The period of the oscillation is 42 months and the exponential damping time scale is 14 months. The POP explains 25% of the variance in the complete observation space and its lag 1 prediction explains 5%. The POP trajectory in EOF space is shown in Fig. 5. Ten trajectories of the unforced system predicted by the POP model are shown. As expected theoretically, all trajectories spiral into a fix point at the origin of the state space.

In a second more general approach we fitted a *nonlinear* model to the data. A two layered feedforward neural network was used as nonlinear function approximator.

If the number n of dimensions of the state space is fixed (four in our case) the neural network model described in section 3 has one free parameter left, the number of hidden neurons N_1. The parameters n and N_1 determine the number of weights of the network, i. e. the total number of model parameters.

By increasing the number of parameters, any model can be fitted to a given data set with an arbitrary small error, but the results become useless when the model repro-

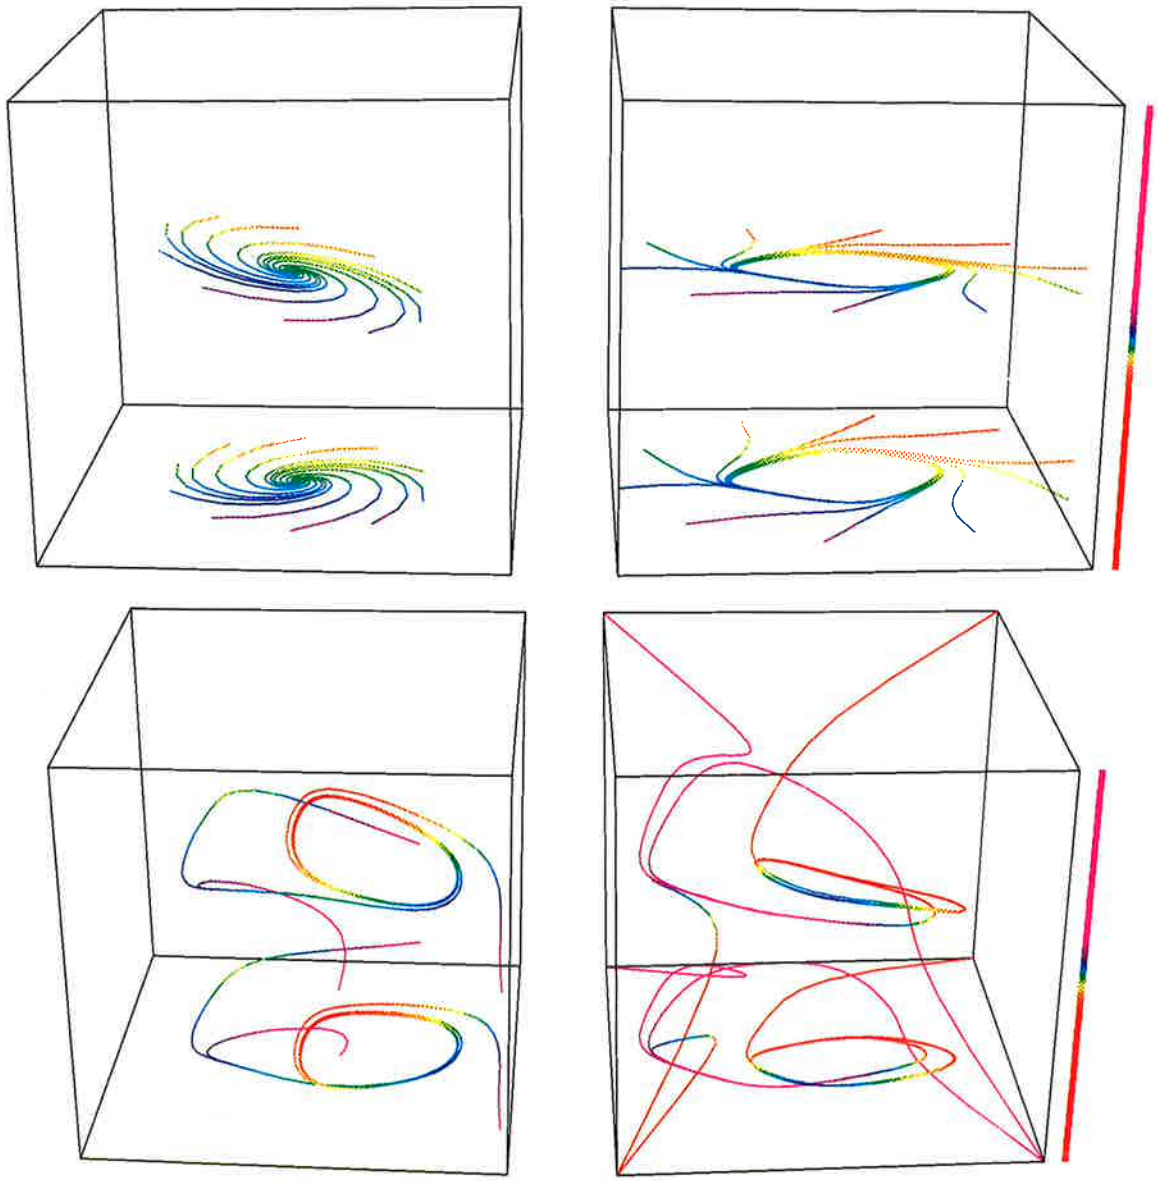

Figure 5: Trajectories in the four-dimensional state space. The spatial dimensions represent the first three EOFs, the color corresponds to the fourth EOF. The color scale with respect to the space scale is displayed at the right margin. All figures show phase space trajectories of the undriven system as predicted by estimated system functions. *Top left:* POP model. All trajectories end up on a fix point attractor. *Top right:* Neural network model with two hidden neurons. The limit cycle attractor is embedded in a plane. *Bottom left:* Neural network model with three hidden neurons. The attractor is embedded in a three-dimensional subspace. Three trajectories starting on points of this subspace are shown. *Bottom right:* Neural network model with four hidden neurons. The attractor is now embedded in the complete four-dimensional state space. Four trajectories are shown, starting in some corners of the four-dimensional hypercube. (In all figures the curve at the bottom of the box is just a projection—a "shadow"—to elucidate the three dimensional structure.)

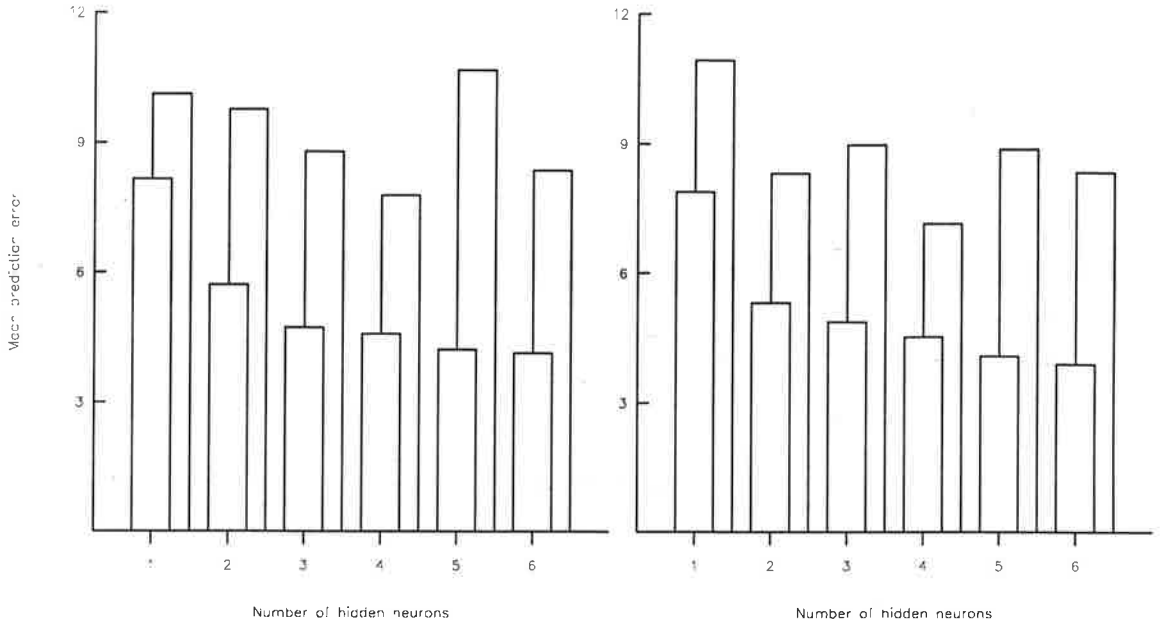

Figure 6: The lag 1 prediction error, i. e. the rms-difference of the observed $\mathbf{x}(t+1)$ and the value calculated from $\mathbf{x}(t)$ by the neural network model as a function of the number of hidden neurons. *Left:* Training with first half of data. *Right:* Training with second half. *Foreground bars:* Error for used half. *Backround bars:* Error for unused half.

duces the stochastic component of the data. To determine the optimal number of hidden neurons, we trained various networks using only one half of the data and checked the lag 1 prediction error for the other independent half. The results are shown in Fig. 6. Using either the first or the second half of the data to train the model, prediction error for the complementary data set was a minimum for a number of four hidden neurons. Larger networks with five or more hidden neurons therefore have too many parameters and reproduce in part the stochastic forcing.

For $N_1 = 1$ (one hidden neuron) the output of the network is of the form

$$y_i^{(2)} = W_{i0}^{(2)} + W_{i1}^{(2)} y_1^{(1)}, \tag{8}$$

i. e. all possible output vectors lie on a straight line in the state space. This prohibits

oscillatory behaviour and restricts the model to fix point attractors. For two, three, and four hidden neurons, the results of training with the complete data set—represented by some typical model state space trajectories—are shown in Fig. 5.

A number of two hidden neurons corresponds to 22 model parameters (section 3). The resulting flow field in the state space converges to a limit cycle attractor, see Fig. 5, top right. This result is qualitatively different from the linear model obtained with the POP analysis, which can only produce damped oscillations converging to a fix point attractor. The unforced temporal evolution of the four EOF coefficients is shown in Fig. 7. Most of the energy is contained in the first two EOFs, i. e. the limit cycle plane is the plane spanned by the first two EOFs. The two coefficient time series are approximately in quadrature. The period of the limit cycle is about 56 months and thus somewhat larger than the POP period of 42 months.

The state space trajectories resulting from a system function represented by a neural network with three hidden neurons, a number which corresponds to 31 model parameters, are presented in Fig. 5, bottom left. The dimension of the output space is increased by one compared with the two-hidden-neuron model. Thus the attractor is now embedded in a three-dimensional subspace of the four-dimensional state space. Although the position and orientation of the attractor has slightly changed, we recover again a limit cycle which has a period of about 44 months (Fig. 7). The variance in direction of the second EOF is smaller than for the two hidden-neuron-model, whereas the variances in direction of the third and fourth EOFs are larger.

The model with the best prediction skill for unused data, namely the neural network with four hidden neurons, recovers also a limit cycle attractor, see Fig. 5, bottom right.

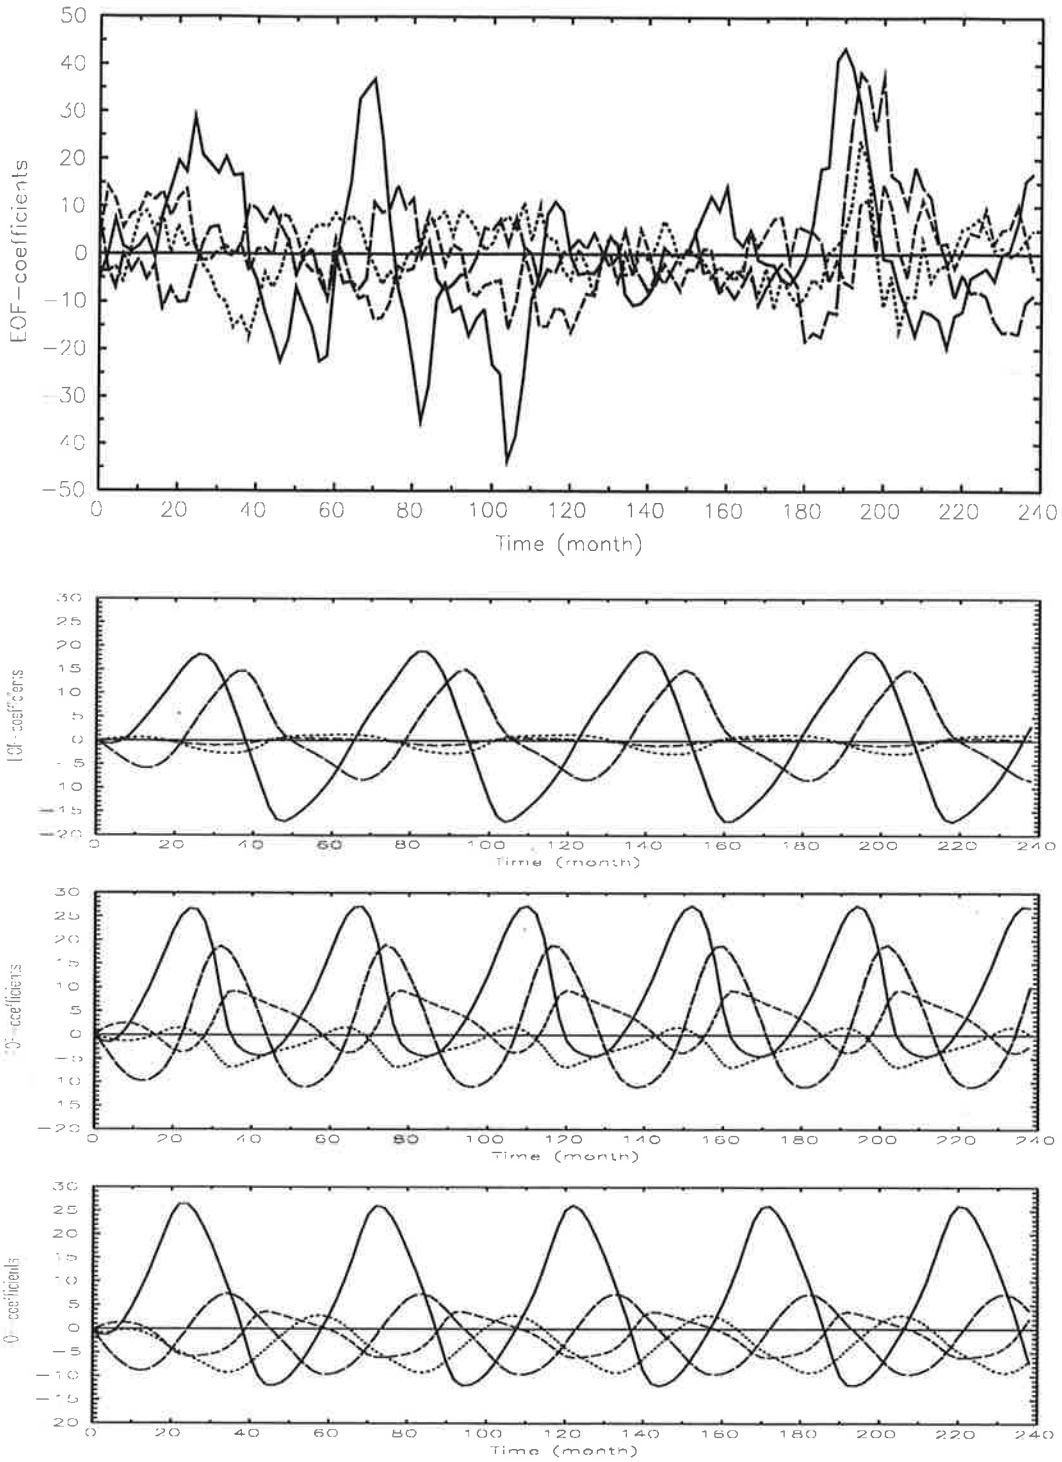

Figure 7: The temporal evolution of the four EOF coefficients as observed (top) and as predicted by the neural network model with two (second from top), three (third from top), and four hidden neurons (bottom). Solid: First EOF. Long, medium, and short dashed: Second, third, and fourth EOFs, respectively.

The period is 49 months, compare Fig. 7. The variance in the direction of the first EOF is more dominant than for the other two models. Here we have 40 model parameters and the possible output vectors fill the complete four-dimensional state space. So the attractor is now embedded in this four-dimensional space. The attractor itself, the limit cycle, has a dimension of one in all cases. The two-hidden-neuron model forces it to be embedded in a plane, while the three- and four-hidden-neuron models allow for a curvature in a third and a fourth dimension, respectively. The qualitative stability of the resulting limit cycle attractor is demonstrated by the similarity of all three network models. The ability of the model to distinguish between fix point and limit cycle attractors is tested in appendix A.

The lag 1 prediction error of the POP model and the neural network models with $N_1 = 1 \ldots 4$ respectively is shown in Fig. 8 and compared with the error of the climatology prediction (the prediction of the mean, i. e. the null-prediction for the zero mean EOF coefficients) and the persistence prediction. While the error of the POP model is slightly larger than that of the persistence for the lag 1 prediction, the error of the neural network model with four hidden neurons is smaller. The lag 1 prediction of the neural network model with four hidden neurons explains 79% of the variance in the four-dimensional EOF space, whereas the lag 1 prediction of the POP model explains 62%.

The neural network model can also be used to make multi-lag predictions of the expected state space trajectory of the system. We make predictions with all three neural network models (two, three, and four hidden neurons) which were trained with the complete available dataset from 1967 to 1986. The last state—November/December 1986—is taken as starting point. It is introduced into the neural network and the output—the state one time step later—is used as new input. This is repeated to produce a predicted

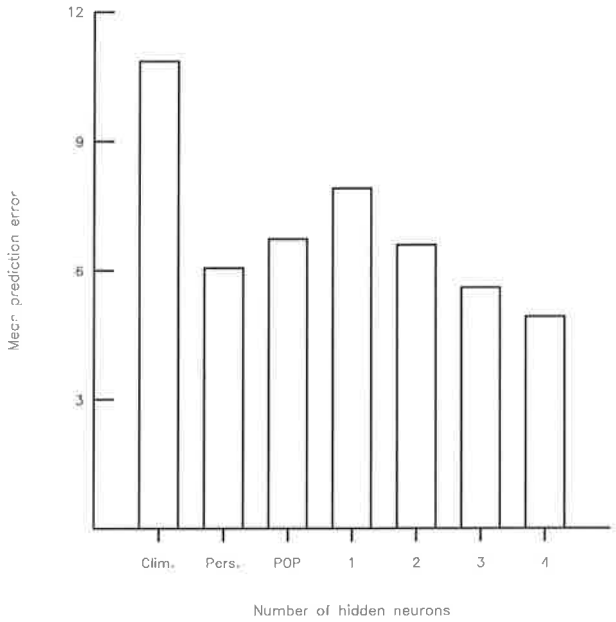

Figure 8: The lag 1 prediction error, i. e. the rms-difference between the observed $\mathbf{x}(t+1)$ and the value calculated from $\mathbf{x}(t)$ by the different models, the POP model and the neural network models $1 \ldots 4$. Training was done with the complete data set. For comparison the error of the climatology prediction—i. e. the prediction of the mean—and the persistence prediction are also shown.

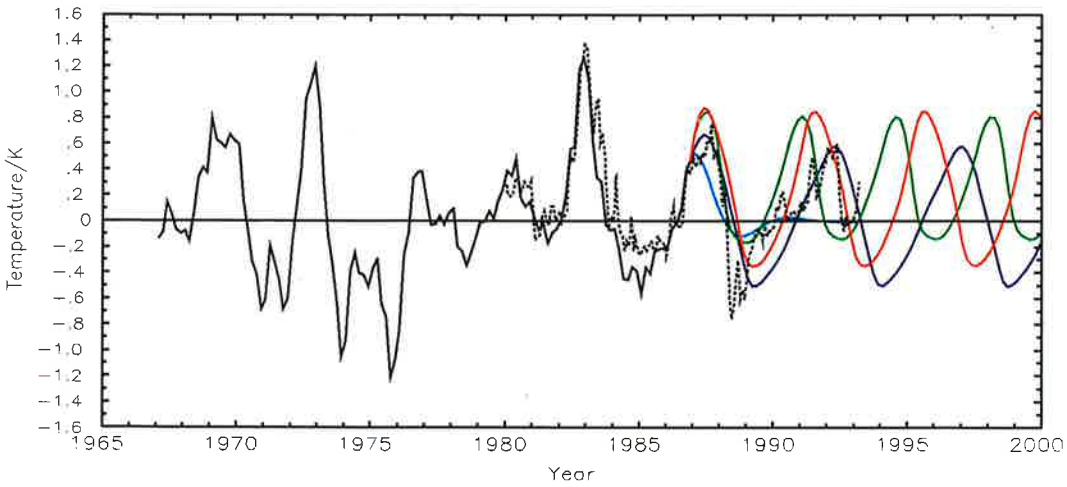

Figure 9: Observed time series of spacially averaged sea surface temperature of the eastern equatorial Pacific obtained from the first four EOFs of the data (solid black line), predictions obtained from the neural network model with two (dark blue), three (green), and four hidden neurons (red), the POP prediction (light blue) and the NINO 3 index (dotted line). The NINO 3 index is normalized to the variance of the temperature observations.

state space trajectory. In Fig. 9 the observed time series of spatially averaged sea surface temperatures of the eastern equatorial Pacific is shown, together with the POP model prediction and the neural network predictions. Due to the different limit cycle periods, the neural network models diverge with increasing lead time.

To compare the prediction with the observed SST since 1986 we use the NINO 3 index of the CAC (Climate Analysis Centre) which was available for the period from January 1980 to March 1993. Predictions and observations are in surprisingly good agreement, especially for the four-hidden-neurons model. All neural network models are considerable superior to the POP model. Note that there is no artificial prediction skill. The data used for the estimation of the models and the initialization of the predictions were available only until 1986.

5 Conclusions

From the observed data, the described neural network recovers a limit cycle attractor. Starting from their physical ocean-atmosphere model, Battisti and Hirst (1989) also obtained a system with a limit cycle attractor, when they reduced their model to one dynamical equation only, the "delayed action oscillator" equation (Schopf and Suarez (1988)). The "delayed action oscillator" scenario, however, is of limited value in understanding the ENSO dynamics due to its ad hoc nature and since it does not pose an eigenvalue problem. A more thorough theoretical analysis of the ENSO dynamics can be found in Münnich *et al.* (1991) and Neelin and Jin (1993). According to these studies, ENSO is largely determined by the first bifurcation from the climatic state leading to a limit cycle. Thus, many aspects of ENSO can be addressed by understanding the linear problem and investigating the most unstable mode of the system linearized about the climatic state. Our results are consistent with this view and indicate that the real system is located above the bifurcation point in parameter space supporting self-sustained oscillations, i. e. the essential dynamics of the El Niño phenomenon is represented by a limit cycle attractor. Additionally, there is some unresolved variability which has to be taken as a kind of stochastic forcing at this point of the investigation. It it possible that this "stochastic" component is in fact also deterministic, so that the complete variability can be described by a strange attractor. But to resolve such an attractor, much more data would be needed. With our limited amount of data only a blurred picture of the attractor can be obtained, and this picture shows a limit cycle attractor.

It has been shown that neural networks with sigmoid nonlinearities (i. e. smoothed

step functions like tanh) are appropriate for approximating the behaviour of dynamical systems. Even with a small number of neurons they can reproduce realistic system behaviour, e. g. damped oscillation on a fix point or limit cycle attractors.

It is often assumed that nonlinear models—which have a comparatively large number of parameters—can only be estimated if very long time series—i. e. a large amount of data—is available. The time series we used is very short (in terms of the time scale of the dynamics) and very noisy. Nevertheless the nonlinear neural network model appears to be significantly superior to the linear POP model in terms of the prediction skill.

A Test of the neural network model

The POP analysis leads always to damped oscillations onto one fix point attractor. This is due to the restrictions of a linear system function. Can the limit cycle attractor recovered by the neural network model be an artefact in the same sense, i. e. does this method always create limit cycles regardless of the dynamics of the data?

To test this we create an artifical data set which is *not* based on a limit cycle attractor. We take the POP which was estimated in section 4 and is shown in Fig. 5 and drive it with stochastic forcing. The strength of the forcing is calculated from the observed data in the following way. For each observed time point we make a lag 1 prediction with the POP model and compare the result with the true state at the next time step. The squared deviation is averaged over all time points to give the strength of the forcing, i. e. we calculate the forcing which would have been necessary to produce the observed data variance if the POP dynamics was true. The driving of the POP with simulated

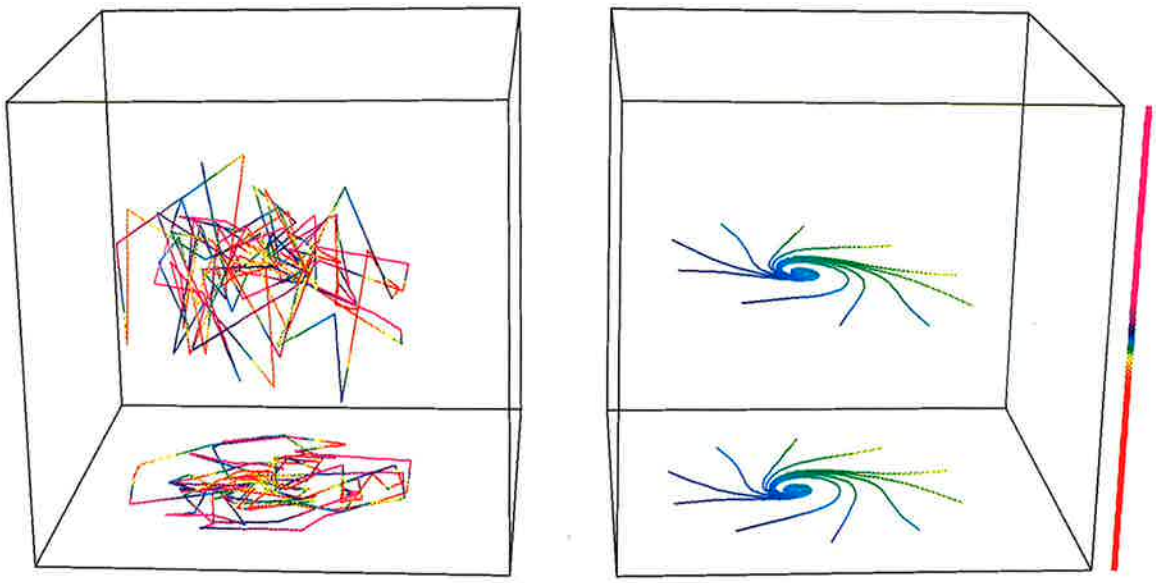

Figure 10: Trajectories in the four-dimensional state space. The spatial dimensions represent the first three EOFs, the color corresponds to the fourth EOF. The color scale with respect to the space scale is displayed at the right margin. *Left:* Simulated phase space trajectory of a stochastically driven POP. *Right:* Phase space trajectories of the undriven system as predicted by the system function estimated with the neural network. All trajectories end up on a fix point attractor. The dynamics is somewhat similar to the underlying POP. (In all figures the curve at the bottom of the box is just a projection—a "shadow"—to elucidate the three dimensional structure.)

stochastic forcing yields the phase space trajectory shown in Fig. 10, left side.

A neural network with two hidden neurons is now trained to reproduce the artificial data. The resultant model dynamics is presented in Fig. 10, right side. It clearly shows the qualitative characteristics of the underlying POP and not a limit cycle attractor. The similarity does not extend beyond the spiral-like attraction onto a fix point. Differences must be expected because the synthetic time series is short (120 bimonth, as the observed

data) and because the number of parameters of the network model is larger than the number of parameters of the POP dynamics.

References

Battisti, D. S., 1988: "The dynamics and thermodynamics of a warming event in a coupled tropical atmosphere/ocean model", *J. atmos. Sci.* **45**, 2889–2919

Battisti, D. S., Hirst, A. C., 1989: "Interannual Variability in a Tropical Atmosphere-Ocean Model: Influence of the Basic State, Ocean Geometrie and Nonlinearity", *J. atmos. Sci.* **46**, 1687–1712

Cane, M. A., Zebiak, S. E., Dolan S. C., 1986: "Experimental forecasts of El Niño", *Nature* **321**, 827–832

Cane, M. A., Münnich, M., Zebiak, S. E., 1990: "A study of self-excited oscillations of the tropical ocean-atmosphere system. Part I: Linear analysis", *J. atmos. Sci.* **47**, 1562–1577

Chao, Y., Philander, S. G. H., 1991: "On the structure of the Southern Oscillation and evaluation of coupled ocean-atmosphere models", TOGA notes no. 3 (April), 1–8

Elsner, J. B., Tsonis, A. A., 1992: "Nonlinear Prediction, Chaos, and Noise", *Bull. Am. met. Soc.* **73**, 49–60

Graham, N. E., White, W. B., 1988: "The El Niño cycle: A natural oscillator of the Pacific Ocean-Atmosphere system", *Science* **240**, 1293–1302

Hasselmann, K., 1988: "PIPs and POPs: The Reduction of Complex Dynamical Systems Using Principal Interaction and Oscillation Patterns", *J. geophys. Res.* **93**,

11,015–11,021

Latif, M., Sterl, A., Maier-Reimer, E., Junge, M. M., 1993: "Structure and predictability of the El Niño/Southern Oscillation phenomenon in a coupled ocean-atmosphere general circulation model", *J. Climate*, in ~~press~~ **6, 701-708.**

Münnich, M., Cane, M. A., Zebiak, S. E., 1991: "A study of self-excited oscillations in a tropical ocean-atmosphere system. Part II: Nonlinear cases", *J. Atmos. Sci.* **48**, 1238–1248

Neelin, J. D., Jin F. F., 1993: "Modes of interannual tropical ocean-atmosphere interactions—a unified view. Part II: Analytical results in the weak coupling limit", *J. atmos. Sci.*, in press

Philander, S. G. H., 1990: "A review of simulations of the Southern Oscillation", International TOGA Scientific Conference Proceedings (Honolulu, Hawaii, U. S. A., 16–20 July 1990), World Climate Research Programme, WCRP-43, WMO/TD-No. 379

Rasmusson, E. M., Carpenter, T. H., 1982: "Variations in tropical sea surface temperature and surface wind fields associated with the Southern Oscillation/El Niño", *Mon. Wea. Rev.* **110**, 354–384

Richman, M. B., Angel, J. R., Gong, X., 1992: "Determination of Dimensionality in Eigenanalysis", Preprint

Rumelhart, D. G., Hinton, G. E., Williams, R. J., 1986: "Learning representations by back-propagating errors", *Nature* **323**, 533–536

Schopf, P. S., Suarez, M. J., 1988: "Vacillations in a coupled ocean-atmosphere model", *J. atmos. Sci.* **45**, 549–566

Storch, H. von, Bruns, T., Fischer-Bruns, I., Hasselmann, K., 1988: "Principal Oscillation Pattern Analysis of the 30- to 60-Day Oscillation in General Circulation Model Equatorial Troposphere", *J. geophys. Res.* **93**, 11,022–11,036

Storch, H. von, Weese, U., Xu, J. S., 1990: "Simultaneous Analysis of Space-Time Variability: Principal Oscillation and Principal Interaction Patterns with Applications to the Southern Oscillation", *Z. Meteorol.* **40**, 99–103

Tang, B., 1992: "El Niño Forecast Using Neural Networks", Preprint

Widrow B., Lehr M. A., 1990: "30 Years of Adaptive Neural Networks: Perceptron, Madaline, and Backpropagation", *Proc. IEEE* **78**, 1415–1441

Zebiak, S. E., Cane, M. A., 1987: "A model El Niño-Southern Oscillation", *Mon. Wea. Rev.* **115**, 2262–2278