

R. Knutti · T. F. Stocker · F. Joos · G.-K. Plattner

Probabilistic climate change projections using neural networks

Received: 30 October 2002 / Accepted: 16 May 2003 / Published online: 5 August 2003
© Springer-Verlag 2003

Abstract Anticipated future warming of the climate system increases the need for accurate climate projections. A central problem are the large uncertainties associated with these model projections, and that uncertainty estimates are often based on expert judgment rather than objective quantitative methods. Further, important climate model parameters are still given as poorly constrained ranges that are partly inconsistent with the observed warming during the industrial period. Here we present a neural network based climate model substitute that increases the efficiency of large climate model ensembles by at least an order of magnitude. Using the observed surface warming over the industrial period and estimates of global ocean heat uptake as constraints for the ensemble, this method estimates ranges for climate sensitivity and radiative forcing that are consistent with observations. In particular, negative values for the uncertain indirect aerosol forcing exceeding -1.2 W m^{-2} can be excluded with high confidence. A parameterization to account for the uncertainty in the future carbon cycle is introduced, derived separately from a carbon cycle model. This allows us to quantify the effect of the feedback between oceanic and terrestrial carbon uptake and global warming on global temperature projections. Finally, probability density functions for the surface warming until year 2100 for two illustrative emission scenarios are calculated, taking into account uncertainties in the carbon cycle, radiative forcing, climate sensitivity, model parameters and the observed temperature records. We find that warming exceeds the surface warming range projected by IPCC for almost half of the ensemble members. Projection uncertainties are only consistent with IPCC if a

model-derived upper limit of about 5 K is assumed for climate sensitivity.

1 Introduction

There is strong observational evidence for a significant warming of the Earth's climate system, from both instrumental records of atmospheric temperature over the last 150 years (Jones et al. 1999), and a recent reanalysis of ocean temperature data of the last 50 years (Levitus et al. 2000). Most of the global-mean temperature increase of the last 50 years can be attributed to anthropogenic influence through the emissions of greenhouse gases and other radiatively active gases (IPCC 2001). Optimal fingerprint methods used for detection and attribution of climate change have shown that the observed atmospheric warming can neither be explained by natural variability, nor by an anthropogenically forced increase in greenhouse gases alone, but is best matched when taking into account different anthropogenic forcings (greenhouse gases and sulfate aerosols) together with natural forcings (variations in the solar irradiance and stratospheric volcanic aerosols) (Santer et al. 1996; Tett et al. 1999; Stott et al. 2000, 2001; Hegerl et al. 2000). Similarly, the observed changes in the oceanic heat content are consistent with those expected from anthropogenically forced model simulations, and natural variability can be excluded as the origin of the observed long-term oceanic warming with very high confidence (Barnett et al. 2001; Reichert et al. 2002).

Future emissions of greenhouse gases will further amplify the human impact on the climate system and therefore increase the need for accurate climate projections based on complex numerical models. A prime concern about such future projections are their rather large uncertainties, which arise from uncertainties in future emissions of radiatively active trace gases and

R. Knutti (✉) · T. F. Stocker · F. Joos · G.-K. Plattner
Climate and Environmental Physics, Physics Institute,
University of Bern, Sidlerstrasse 5,
3012 Bern, Switzerland
E-mail: knutti@climate.unibe.ch

aerosols, in converting emissions to concentration changes, in calculating the radiative forcings from the increased atmospheric concentrations, and in estimating climate changes in response to forcing changes.

A large and usually neglected uncertainty in climate change projections arises from the uncertainties in aerosol forcing. Anthropogenic emissions of aerosols affect the radiative balance of the atmosphere in several ways. The direct effects include the extinction of sunlight by aerosols as well as impacts of aerosol absorption on the temperature and humidity profiles, which affect cloud cover and formation. Furthermore, there are indirect aerosol effects. The first indirect effect relates to the increase in cloud optical depth and cloud albedo due to an increase in the number and a decrease in the size of cloud droplets. The second indirect effect refers to an increase in the cloud liquid water content, cloud height or cloud lifetime due to a reduced precipitation efficiency because of the decrease in the cloud droplet size.

Great efforts have been made in the Third Assessment Report (TAR) of the Intergovernmental Panel on Climate Change (IPCC 2001) to objectively quantify the uncertainties and probabilities of the main conclusions. However, a satisfactory treatment of the uncertainties for several key findings was impossible. In particular, the uncertainties of the radiative forcing and of the climate sensitivity could only be estimated as a range covered by several studies using different models. Furthermore, the projections of future climate changes lack an objective uncertainty estimate, both because there are no probabilities attached to different scenarios, and because the uncertainties for a specific scenario are only given as ranges covered by different climate models. These ranges cannot be interpreted statistically, i.e. both the distribution of probability within the range and the probability for a projection being outside the range are unknown. Since the number of simulations that can be performed using comprehensive climate models is strongly limited due to their computational cost, a rigorous analysis of the projection of uncertainties from comprehensive models currently remains unfeasible.

The observed surface warming and ocean heat uptake can be used to constrain parameters in the climate system as well as in projections of future climate. Two partly complementary methods have been used so far for quantitative uncertainty analysis: optimal fingerprint and Monte Carlo methods. Optimal fingerprint patterns detected in the observed warming and in global warming simulations from atmosphere–ocean general circulation models (AOGCMs) can be used to estimate a constant scaling factor which is used to scale future climate projections, if they over- or underestimate the observed warming. Climate variability from a constant climate control integration is then used to calculate the uncertainty of this scaling factor, resulting in a probability range for the projected future warming (Allen et al. 2000, 2002; Stott and Kettleborough 2002). Monte

Carlo methods have been used to calculate large model ensembles, taking into account the uncertainties in model parameters and input data. After selecting those simulations that are consistent with the observed warming, probability density functions for model parameters as well as for the projected warming can be calculated (Wigley and Raper 2001; Forest et al. 2002; Knutti et al. 2002). In the ensemble approach, uncertainties of any type can be properly considered. The main drawback is the fact that only simple models are efficient enough to calculate large ensembles. Natural variability, regional information and certain complex feedbacks are therefore neglected due to the limited spatial and temporal resolution and missing atmospheric dynamics.

Although these simple climate models are efficient, large ensembles are still limited by CPU time. Thus we apply here the concept of neural networks to approximate the climate response function in order to increase the efficiency of the ensemble method. Neural networks are an attempt to mathematically describe the functioning of neurons and therefore to model some features of the human brain. The purpose of this study is to extend the ensemble approach by Knutti et al. (2002) by introducing a neural network climate model substitute. The performance of neural networks in approximating functions is used here to predict a climate response from selected input parameters, without explicitly integrating the climate model. This dramatically improves the efficiency of the ensemble method, allowing sensitivity studies and a more detailed analysis of the processes contributing to the total uncertainty. Probabilistic results are presented for future steric sea level rise and for the expected reduction of the Atlantic thermohaline circulation. Furthermore, feedback mechanisms related to changes in the oceanic and biospheric carbon cycles due to a warming climate are investigated.

This study is organized as follows. The climate model, the ensemble procedure and the design of the neural network are described in Sect. 2. The resulting constraints on climate sensitivity, radiative forcing, future warming, the weakening of the Atlantic thermohaline circulation and on steric sea level rise are discussed in Sect. 3. Discussions and conclusions follow in Sect. 4. Appendix 1 explains the details of the carbon-cycle climate feedback. A short introduction to neural networks, and the specific problems in constructing the climate model substitute are given in Appendix 2.

2 Model and methods

2.1 Climate model

The climate model of reduced complexity used in this study consists of a zonally averaged dynamical ocean model (Stocker and Wright 1991; Wright and Stocker 1991) resolving the Atlantic, Pacific, Indian and Southern oceans, coupled to a zonally and vertically averaged energy- and moisture-balance model of the atmosphere

(Stocker et al. 1992; Schmittner and Stocker 1999). The annual mean version is used for efficiency, since differences between it and the seasonal version (Schmittner and Stocker 2001) are negligible on the spatial and temporal scales considered. The simple geographical structure restricts studies to scales of thousands of kilometres, on which results are generally consistent with those of comprehensive 3D models. While the climate sensitivity of comprehensive models is determined by the strength of the resolved feedback mechanisms, the energy-balance atmosphere requires a different approach. We specify the additional radiative forcing at the top-of-the-atmosphere as $\Delta F_{\text{toa}}(t) = \Delta F_{\text{dir}}(t) + \lambda \cdot \Delta T_{\text{atm}}(t)$, where ΔF_{dir} is the direct radiative forcing reconstructed over the industrial period and prescribed for the future according to the assumed scenario. Feedback processes which increase the climate sensitivity are represented by the feedback term $\lambda \cdot \Delta T_{\text{atm}}(t)$, where ΔT_{atm} is the time-dependent atmospheric temperature increase and λ is a constant, which is prescribed to obtain different climate sensitivities for the same radiative forcing. The climate sensitivity for each model simulation is determined by stabilizing the direct radiative forcing ΔF_{dir} at 3.71 W m^{-2} , corresponding to a doubling of the preindustrial atmospheric CO_2 concentration (Myhre et al. 1998), and diagnosing the equilibrium atmospheric temperature increase after 3000 years of integration.

The anthropogenic radiative forcing from changes in well-mixed greenhouse gases (CO_2 , CH_4 , N_2O , SF_6 and 28 halocarbons including those controlled by the Montreal Protocol), stratospheric and tropospheric O_3 , the direct forcing of black and organic carbon and sulfate, stratospheric H_2O due to CH_4 changes, and the indirect effects of aerosols are individually prescribed from reconstructions for the year 1765 to 2000, follow a SRES scenario (Nakićenović et al. 2000) from year 2000 to 2100 and are then slowly adjusted to a constant value of 3.71 W m^{-2} (corresponding to doubling preindustrial atmospheric CO_2) in order to diagnose the climate sensitivity. The radiative forcing is calculated from the changes in concentration using simplified expressions as described in detail by Joos et al. (2001). The radiative forcing by volcanoes and variations in solar irradiance are prescribed for the historical period (Crowley 2000). Albedo changes due to land use changes and radiative forcing by dust are not considered. Details of the radiative forcing components considered are given in Table 1.

Table 1 Radiative forcing components considered in this study

Forcing	Best guess	Uncertainty	Distribution
CO_2 , CH_4 , N_2O , SF_6 , halocarbons	2.43 W m^{-2}	5%	Normal
Stratospheric O_3	-0.15 W m^{-2}	33%	Normal
Tropospheric O_3	0.38 W m^{-2}	22%	Normal
Sulfate	-0.40 W m^{-2}	Factor of 2	Log-normal
Indirect aerosol effects	NA	-2 to 0 W m^{-2}	Uniform
Organic and black carbon	-0.10 W m^{-2}	Factor of 3	Log-normal
Stratospheric water vapour	0.03 W m^{-2}	Factor of 3	Log-normal
Volcanic	NA	Factor of 2	Log-normal
Solar	0.32 W m^{-2}	33%	Normal

The best guess values indicate the radiative forcing at year 2000 relative to 1750. For the standard case, a Gaussian uncertainty distribution is assumed where the uncertainties are given in percent, and a log-normal distribution is used where the uncertainty is given as a factor. The uncertainties indicated are derived from the assumption that the uncertainties given by IPCC represent a ± 2 standard deviation range. A range with equal probability is assumed for the indirect aerosol forcing, since no best guess is available. A best-guess value for the volcanic forcing is meaningless, since it varies very strongly with time, but is almost zero for both the years 1750 and 2000

2.2 Carbon cycle climate feedback

In a previous study (Knutti et al. 2002), we neglected the uncertainty in converting future greenhouse gas emissions given by the SRES scenarios into atmospheric concentrations. This is improved here by introducing a “carbon cycle-climate feedback parameterization” with an associated uncertainty. For a warmer climate, oceanic and terrestrial uptake of anthropogenically emitted carbon is reduced due to a decreased solubility of CO_2 in warmer surface water, (e.g. Sarmiento and Le Quéré 1996; Joos et al. 1999; Plattner et al. 2001) and due to increased soil respiration and forest dieback (Cox et al. 2000; Friedlingstein et al. 2000; Joos et al. 2001). Both effects increase atmospheric CO_2 compared to a situation with constant climate. The difference in projected CO_2 between model simulations with global warming and without global warming quantifies the feedback between climate change and the global carbon cycle. It is expressed here as the ratio γ of the change in radiative forcing per degree surface warming. Radiative forcing from atmospheric CO_2 , $RF_{\text{CO}_2}(t, \Delta T_{2000})$ is thus calculated for each scenario (A1, B2) as a function of the simulated global mean surface temperature change since year 2000, $\Delta T_{2000}(t)$ and a baseline trajectory of future CO_2 forcing, $RF_{\text{CO}_2}(t, \Delta T_{2000} = 0)$:

$$RF_{\text{CO}_2}(t, \Delta T_{2000}) = RF_{\text{CO}_2}(t, \Delta T_{2000} = 0) + \gamma \cdot \Delta T_{2000}(t) \quad (1)$$

In the Monte Carlo simulations, a best estimate for γ of $0.25 \text{ W m}^{-2}/\text{K}$ and a 1-sigma uncertainty of $0.05 \text{ W m}^{-2}/\text{K}$ have been applied. The value of the feedback factor γ and the baseline for CO_2 radiative forcing were derived separately from simulations with the Bern Carbon Cycle-Climate (Bern CC) model (Joos et al. 2001; Gerber et al. 2003). Within one simulation, the feedback parameter γ is kept constant in time. Details of this feedback parameterization are given in Appendix 1.

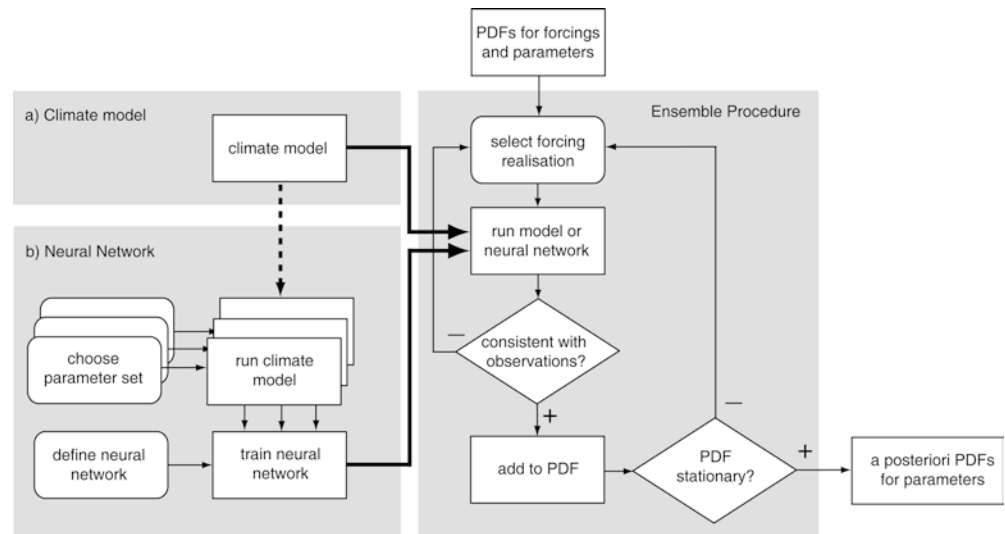
2.3 Ensemble procedure

The ensemble method is essentially based on a Bayesian approach that calculates a posteriori probabilities for uncertainties and parameters in the climate system given the observational evidence. The procedure is outlined in Fig. 1a and starts by defining a priori assumptions concerning the probability distributions of model parameters and input uncertainties R_1 to R_N . The components R_i are assumed to be independent. Since the radiative forcing ranges are derived from individual studies, and climate sensitivity is derived from models, this is a reasonable assumption. A large number (10^3 to 10^4) of simulations is then explicitly calculated using the climate model and changing the values of R_i in each simulation randomly. Probability density functions (PDFs) are calculated afterwards from those simulations which are consistent with the observed surface warming and ocean heat uptake (right part of Fig. 1). In this study, the values for R_i are chosen randomly in a way that the probability distribution of each parameter R_i matches the corresponding a priori uncertainty distribution assumed for R_i . While sampling the values for R_i randomly is probably less efficient overall than sampling in a systematic way, it allows one to monitor the emerging results and to increase the accuracy continuously until the desired ensemble size is reached.

The main uncertainties affecting climate change projections are the uncertainties in the radiative forcing, in climate sensitivity, in the carbon cycle feedback, and in climate model parameters. All of these uncertainties are addressed in this study using the ensemble method. The uncertainty in the carbon cycle, that has been neglected in an earlier study (Knutti et al. 2002), is taken into account here by assuming a Gaussian distribution with a mean value of $0.25 \text{ W m}^{-2}/\text{K}$ and a standard deviation of $0.05 \text{ W m}^{-2}/\text{K}$ for the carbon cycle feedback parameter γ . When assuming these numbers, the two standard deviation range approximately covers the published results for the carbon cycle climate feedback (see climate model description).

The uncertainty in the radiative forcing is considered separately for each forcing component listed in Table 1. Following the approach of Boucher and Haywood (2001), we assume a normal

Fig. 1 Flow chart of the preparation and the ensemble procedures with which the a posteriori probabilities were determined for parameters, radiative forcing components and future warming



distribution for the uncertainty of each forcing component where absolute errors are given by the IPCC, and a log-normal distribution where the errors are given as a factor. The IPCC errors are taken to be two standard deviations for the standard case, identical to Knutti et al. (2002), although the IPCC attaches no statistical meaning to them. For the indirect aerosol forcing, no best guess is available, and any value between 0 and -2 W m^{-2} is given equal probability. The sensitivity of the results to these assumptions is discussed in Sect. 3.2. A factor of two is assumed for the uncertainty of the volcanic forcing (Hansen et al. 1998).

The climate sensitivity is widely used to indicate the magnitude of a climate model temperature response to changes in radiative forcing. The IPCC (2001) has agreed on a probable range of 1.5 to 4.5 K for climate sensitivity, covering the values estimated from most of the current comprehensive 3D models. However, values exceeding 4.5 K cannot be ruled out by the observed warming (Andronova and Schlesinger 2001; Gregory et al. 2002; Knutti et al. 2002). We therefore assume an a priori range of 1 to 10 K for climate sensitivity with equal probability, and consider the observations as the only objective constraint. It has already been shown that the assumption of an upper limit of 4.5 K is critical for the IPCC (2001) conclusions, and reduces the upper limit of the projected warming of the next century considerably (Knutti et al. 2002). The agreement of the atmospheric warming derived in our model with the instrumental record and the consistency of the results from this simplified model (Knutti et al. 2002) with projections obtained with more complex models (Stott and Kettleborough 2002; Zwiers 2002) provides support that our simplified approach of using a feedback parameter to obtain different model climate sensitivities is justified.

Apart from climate sensitivity and radiative forcing, the largest uncertainties in the climate model used here arise from the parameterization of ocean subgrid-scale mixing processes, which strongly influence the strength of the thermohaline circulation and the diffusive mixing, resulting in a large uncertainty in the oceanic heat uptake in a warming climate. To account for that, the number of ocean model versions considered here is increased to ten, compared to the five in an earlier study. The first three model versions are identical to the versions used earlier (Knutti et al. 2000; Knutti and Stocker 2000, 2002), with a vertical diffusivity of $5 \cdot 10^{-5} \text{ m}^2 \text{ s}^{-1}$ and the three different subgrid-scale mixing parameterizations (a) horizontal/vertical diffusion, (b) isopycnal/diapycnal diffusion and (c) isopycnal/diapycnal diffusion including a Gent-McWilliams advective mixing parameterization (Gent and McWilliams 1990; Gent et al. 1995), and have been described in detail by Knutti et al. (2000). A second set of three model versions uses the same three mixing parameterizations, but with parameter settings used in an earlier study with this climate model (Schmittner and Stocker 1999). To account for the large uncertainty in the vertical

diffusivities used in ocean models, a third set of four model versions with standard horizontal/vertical diffusion parameterization and variable vertical diffusivities 2, 4, 6 and $8 \cdot 10^{-5} \text{ m}^2 \text{ s}^{-1}$ is included in the ensemble. This list of model versions is certainly not exhaustive, but, within the framework of simplified ocean models, these choices approximately cover the uncertainties associated with large-scale ocean mixing.

2.4 Neural network substitute

The calculation of probability density functions requires a relatively large number of simulations. However, the computational cost of most complex climate models exceeds the available CPU capacities by orders of magnitude, if such large ensembles are necessary. Ensemble numbers are limited to a few thousand even when using efficient climate models. This excludes extensive sensitivity studies and the investigation of processes under different model assumptions. A promising way to increase the efficiency of the ensemble procedure is the use of a model substitute, based on the following considerations. For every randomly chosen vector $\vec{R} = (R_1, R_2, \dots, R_N)$, whose components represent the uncertainties or forcing scaling factors in the climate model, a response $C(\vec{R}, t)$ is calculated by the climate model. If a function $F(\vec{R}, t)$ exists such that $C(\vec{R}, t) = F(\vec{R}, t) + \varepsilon$, with the error term ε being small compared to C for all \vec{R} and t , then the climate model substitute $F(\vec{R}, t)$ could be used to speed up the ensemble procedure. In other words, to use the efficient fitted function $F(\vec{R}, t)$ instead of explicitly calculating $C(\vec{R}, t)$ for every member of the ensemble, the part of the model response determined by the parameters R_i must be large, and the chaotic part of the response (e.g. representing natural variability or nonlinear or abrupt changes in the THC not predictable from \vec{R}) must be negligible. This condition is satisfied for all model quantities on the time scales considered here, since the energy-moisture-balance atmosphere model exhibits no natural variability. Therefore, a model substitute can provide results of similar quality to those of the original climate model in this case, and it will have the same limitations of not simulating internal variability, which is clearly not negligible in the real world. This however is considered by requiring that only part of the observed warming is explained by radiative forcing.

Pulse response functions and orthogonal polynomials have been successfully used to approximate climate model responses (e.g. Forest et al. 2002; Joos and Bruno 1996). Here we propose a different approach using a neural network. Of the many types of neural networks with specific properties explored so far, the multi-layer feed-forward networks have proven to be particularly efficient in approximating a wide range of functions. The main difficulty arising in practical applications is the task of defining a neural

network with optimal performance at minimal computational cost. A short introduction to neural networks, the discussion of a few general properties and the choice of an appropriate network size and training set are discussed in Appendix 2.

The main advantage of neural networks compared to other methods is that any relationship between input parameters and output data can be approximated. Even thresholds or qualitatively different behaviour in different parts of the parameter space can be modelled. Further, the training simulations can be randomly distributed in the parameter space. For regular sampling methods like the latin hypercube sampling (Forest et al. 2002), it is often difficult to know in advance the number of simulations that is required to get a good approximation to the original model. In contrast to these methods, the number of randomly distributed simulations used to train the neural network can be increased continuously until the performance of the neural network is sufficient.

The neural network-based climate model substitute is two to three orders of magnitude faster than the climate model used here. Even when considering the cost of calculating a training set, the neural network ensemble method is an order of magnitude or two more efficient (depending on the ensemble size and the size of the training set) and also more flexible than calculating the ensemble explicitly with the climate model. For example, the a priori assumptions for the uncertainties can be changed easily to investigate sensitivities, without recalculating the training set or retraining the network.

To summarize, the ensemble procedure including the neural network substitute is sketched in Fig. 1b. First, a priori probabilities for the errors and uncertainties have to be assumed. Second, a set of 1000 simulations is calculated by varying the uncertainties randomly. Third, the neural network is designed and trained using half of the simulations for training and the rest for independent validation. Fourth, the actual ensemble is generated using the neural network to predict the climate response from the uncertainty parameters. Those simulations consistent with observations are selected to calculate an a posteriori probability density function of whatever quantity is desired. The ensemble size is continuously increased until the PDF is stationary. For all the results presented, the initial unconstrained ensemble encompasses at least 10^6 ensemble members generated by the neural network. The probability density functions obtained can therefore be considered as stationary with a high degree of accuracy.

2.5 Model and neural network performance

The performance of the climate model as well as the neural network substitute is shown in Fig. 2. The plotted simulation is randomly chosen for illustration among those simulations which are consistent with the observations. In general, surface warming calculated from the climate models (solid) agrees well with observations (shaded band, data from Jones et al. 1999). Both the trend and most of the decadal variations are reproduced. However, some features like the almost constant temperatures between 1940 and 1970 and the strong warming after 1980 are not well reproduced, indicating that either the external forcing is not entirely correct or that part of the observed warming is due to internal processes. For the ocean heat uptake over the last 40 years, only the trend of the data (Levitus et al. 2000) can be reproduced, very similar to results obtained by comprehensive models (Barnett et al. 2001; Reichert et al. 2002). Even the most complex models are currently not able to simulate such observed internal variability, e.g. caused by the El Niño-Southern Oscillation (Levitus et al. 2001; Barnett et al. 2001; Reichert et al. 2002).

The performance of the neural network in approximating the climate model response is also illustrated by the simulation shown in Fig. 2. The approximated response from the trained neural network (dashed) agrees very well with the explicit climate model simulation (solid). The error of the neural network is negligible compared to the uncertainties related to the model setup, parameters and input data, and the deviations of the neural network are not systematic.

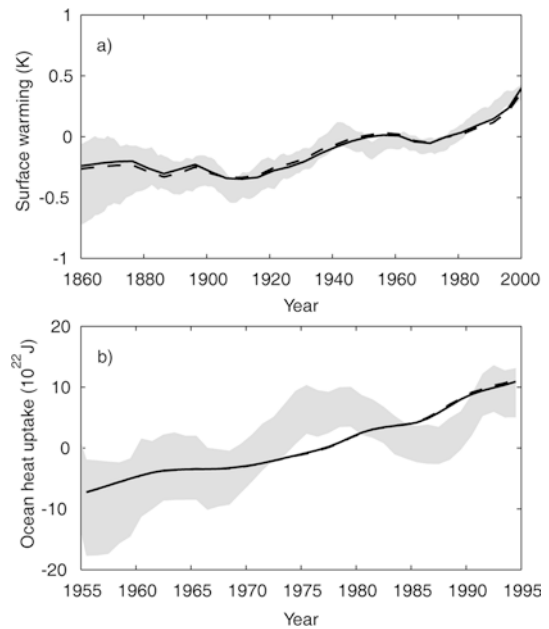


Fig. 2 a Surface warming and b ocean heat uptake simulated by the climate model (solid) and approximated by the trained neural network (dashed) over the observational period. The simulation shown here is randomly chosen for illustration among those matching the observational constraints. One standard deviation of the observations are shown as shaded bands for the reconstructed global mean surface warming (Jones et al. 1999) and for the global ocean heat uptake (Levitus et al. 2000). Note that the time axes are not identical

3 Results

3.1 Climate sensitivity

The climate sensitivity for ocean–atmosphere models is commonly expressed as the equilibrium surface air temperature increase for a doubling of the preindustrial atmospheric CO₂ concentration, and varies considerably between different models (IPCC 2001). For comprehensive coupled atmosphere ocean general circulation models (AOGCM), the climate sensitivity is internally determined by physical, chemical and biological (feedback) processes and by the way they are parameterized in the model. Our incomplete knowledge of the cloud feedbacks in particular contributes to the large uncertainty in climate sensitivity. The most recent assessment by the IPCC (2001) confirmed the range of 1.5 to 4.5 K established in earlier studies (Shine et al. 1995), without indicating any statistical interpretation of that range.

It has recently been suggested that the largely uncertain climate sensitivity can be constrained by relating the reconstructed radiative forcing over the industrial period to the observed surface air warming and the observed ocean heat uptake. The requirement that the modelled warming matches the observed warming should thus place a strong constraint on anthropogenically forced climate models, and the ocean heat uptake should impose an even tighter restriction

than the atmospheric warming because of the ocean's large heat capacity (Barnett et al. 2001). However, these conclusions are hampered by the incompleteness of climate models, by the uncertainty in the temperature records and the uncertainty of the reconstructed radiative forcing over the last 250 years (Andronova and Schlesinger 2001; Gregory et al. 2002; Knutti et al. 2002).

The upper limits for the probability density function for climate sensitivity found in the different studies (Andronova and Schlesinger 2001; Forest et al. 2002) are therefore always larger than the IPCC estimate of 1.5 to 4.5 K derived from models. In our case, when assuming any value from 1 to 10 K to be equally probable, the observational constraints cannot place an upper limit below 10 K for climate sensitivity. Although values around 2.5 K are most probable and consistent with most AOGCMs, large climate sensitivities cannot be ruled out by this method. The IPCC range of 1.5 to 4.5 covers only about 50% of the resulting PDF shown in Fig. 3. A significant reduction in the uncertainties of the observational datasets as well as better constraints on the reconstructed radiative forcing are required to considerably reduce the uncertainty in climate sensitivity and, at the same time, the uncertainties for long-term climate projections.

3.2 Radiative forcing

Climate model simulations including negative radiative forcings by anthropogenic aerosols have achieved much better agreement with historical temperature observations than those considering greenhouse gas forcings only (e.g. Santer et al. 1996; Stott et al. 2000). However, estimates of the net negative forcing by aerosols have uncertainties that are comparable in magnitude to the positive forcing by greenhouse gases. In particular, the indirect forcing component of the aerosols has a very large uncertainty. Thus, confidence in attributing

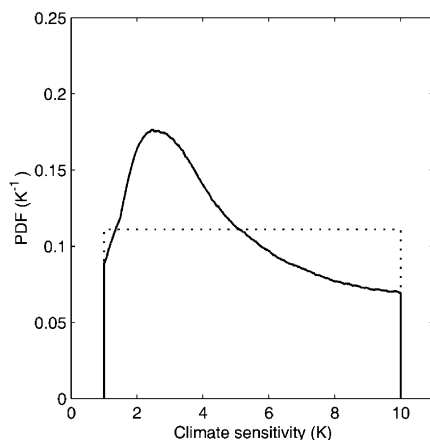


Fig. 3 PDF for climate sensitivity as constrained by the ensemble simulations and observed surface warming and ocean heat uptake. The dotted line represents the initial assumption of the PDF

observed warming directly to forcing by enhanced greenhouse gas concentrations depends upon how well the aerosol forcings are determined.

Two distinct and independent approaches have been used to quantify the indirect aerosol forcing, the most uncertain of all radiative forcing components. First, measured or calculated aerosol properties combined with estimated global distributions from chemical transport models can be used to estimate the aerosol forcing. Second, the strength of the indirect aerosol forcing can be inferred from the constraint that the observed Earth warming of the industrial period places on the total radiative forcing. In general, the latter approach yields smaller ranges than modelling the aerosols explicitly. For the a priori assumption, we adopt the range of 0 to -2 W m^{-2} proposed by the IPCC (2001) for the indirect aerosol forcing. Since no best guess is available, we assume equal probability in the whole range (Fig. 4a, dotted line). When selecting the simulations from the ensemble that are consistent with the observed surface warming and ocean heat uptake, the PDFs shown in Fig. 4a are obtained. The solid line includes all climate sensitivities as constrained from Fig. 3, and yields a probability of 93% for the indirect aerosol effect being smaller than 1.2 W m^{-2} in

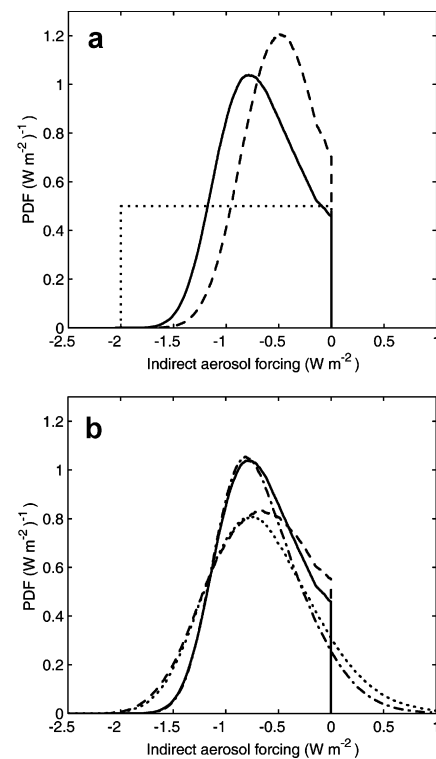


Fig. 4 a PDF for the indirect aerosol forcing and standard case. *Solid*: all climate sensitivities; *dashed*: IPCC climate sensitivities; *dotted*: initial assumptions. **b** PDF for the indirect aerosol forcing and different a priori forcing error assumptions for the IPCC forcing errors and the indirect aerosol forcing (denoted in the form $??$). *Solid* (standard as in **a** $2\sigma/-2$ to 0 flat; *dashed*: $1\sigma/-2$ to 0 flat; *dash-dotted*: $2\sigma/1 \pm 1$ normal; *dotted*: $1\sigma/1 \pm 1$ normal. See main text for explanations

magnitude. If the IPCC upper limit of 4.5 K for climate sensitivity is assumed, the probability increases to 99%. These results are similar to those of Knutti et al. (2002), but not necessarily identical, since more ocean model versions are considered here.

One might argue that these results are sensitive to the a priori choices of the different forcing assumptions. Basically, two assumptions can be relaxed to test this sensitivity. First, we have assumed all IPCC forcing uncertainties to represent one standard deviation instead of two to ensure that the results are similar when using less restrictive a priori assumptions (see Boucher and Haywood 2001, for a discussion of the IPCC radiative forcing uncertainties). The a priori PDFs are therefore stretched by a factor of two compared to the standard case. The assumption for the indirect aerosol effect is left unchanged. Compared to the standard case (Fig. 4b, solid), the resulting PDF (dashed) still excludes values larger than -1.2 W m^{-2} in magnitude for the indirect forcing with a probability of 84%. Second, we have relaxed the a priori assumption for the indirect aerosol effect alone to a normal distribution with the mean at -1 W m^{-2} and a standard deviation of 1 W m^{-2} , covering both values larger than zero and values smaller than -2 W m^{-2} . The changes for strongly negative values are negligible in this case (dash-dotted line). However, positive values for the indirect aerosol forcing cannot be excluded by this method. When relaxing both assumptions together, the two effects combine, but largely negative forcings can still be excluded. For restricted climate sensitivities (1.5 to 4.5 K), the constraints are even stronger. Regarding the possibility for positive values, there is general agreement from chemical transport and radiation models that the indirect aerosol forcing is very unlikely to be positive (IPCC 2001). Summarizing the results, we can therefore constrain the indirect aerosol forcing and forcings not considered in this study to be between -1.2 and 0 W m^{-2} with high confidence, independent of the underlying assumptions on forcing uncertainties and climate sensitivities.

Our results suggest that similar constraints cannot be derived for the other forcing components. In Fig. 5 the uncertainties of all forcing components are expressed as the 5 to 95% probability range of the scaling factor used to vary the corresponding forcing component. The left error bar (squares) always denotes the a priori assumption, while the right error bar (circles) indicates the constraint obtained by the ensemble method and the observations. It is obvious that apart from the indirect aerosol effect, the uncertainty of the forcing components cannot be narrowed. The upper limit assumed for the volcanic forcing is probably too large, but all other assumptions are consistent with the observations.

3.3 The probability for future warming

The updated projections for the future surface warming were possibly the most anticipated result of IPCC TAR

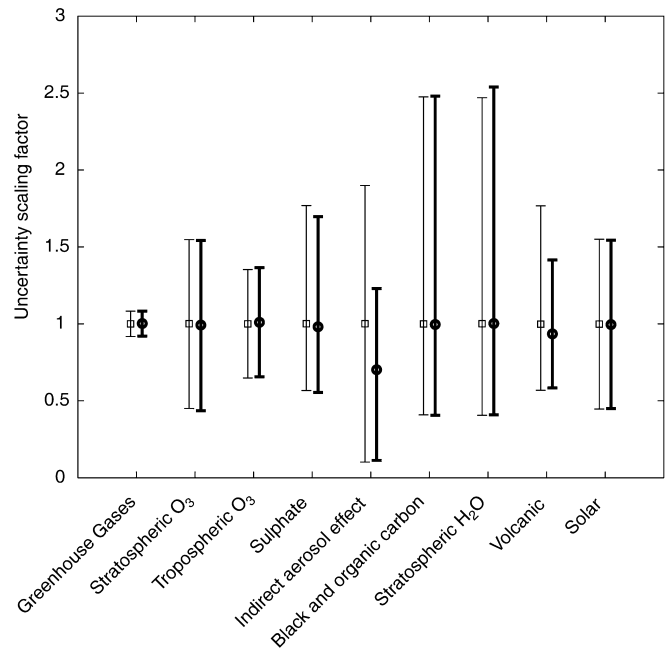


Fig. 5 Median and 5% to 95% ranges of the uncertainty forcing scaling factors for the standard case. The left error bars with squares indicate initial assumptions, the right error bars with circles denote the range as constrained by the ensemble method and the observed surface warming and ocean heat uptake

(IPCC 2001). The warming range of 1.4 to 5.8 K finally presented in the TAR for the year 2100 was significantly wider and shifted towards higher values than in the previous assessment primarily due to lower sulfur emissions in the recently developed IPCC scenarios (Nakićenović et al. 2000) as compared to the earlier used IS92 emission scenarios (Legget et al. 1992). The IPCC (2001) did not attach statistical bounds to their temperature ranges. One therefore cannot decide whether a warming above 5 K is likely to occur or not, and neither do we have any idea about the probability for a warming outside of the specified range. Thus, there is a growing need for climate projections with objectively determined probabilities and uncertainties (e.g. Schneider 2001). The first probabilistic projections based on ensembles without using observational constraints were presented by Wigley and Raper (2001). Other probabilistic projections were obtained by scaling future warming by a linear factor derived from fingerprint patterns (Allen et al. 2000, 2002; Stott and Kettleborough 2002). Here we present results applying the neural network based ensemble method and using the observed surface warming and ocean heat uptake to constrain the model responses.

The uncertainty in the projected temperature arises from two fundamentally different sources. First, the IPCC developed numerous illustrative scenarios with different emission pathways for greenhouse gases and aerosols to account for the uncertainty in future technical, economical and social development (Nakićenović et al. 2000). Second, our incomplete knowledge of the

climate system, the uncertainties in observational datasets, the simplifications made in the climate models and the limited computational resources introduce an uncertainty in the calculated projections for a single scenario. Here, we only address the latter uncertainties by selecting the two illustrative scenarios B1 and A2 (Nakićenović et al. 2000). Since no probabilities have been attached to the individual SRES scenarios by IPCC (2001), we have abstained from calculating an overall probabilistic estimate of future warming.

The observed surface warming and ocean heat uptake can be used to select those ensemble simulations that correctly reproduce the warming over the industrial period, yielding constraints on model parameters and radiative forcing components (see previous sections). The “consistent” set of simulations can now be extended to the year 2100 to calculate probabilities for future climate change. The projected global mean surface warming is shown in Fig. 6 for scenarios B1 (thin lines) and A2 (thick lines). The four panels show the PDFs for the time slices of the years 2025, 2050, 2075 and 2100. The projected warming increases with time, and the

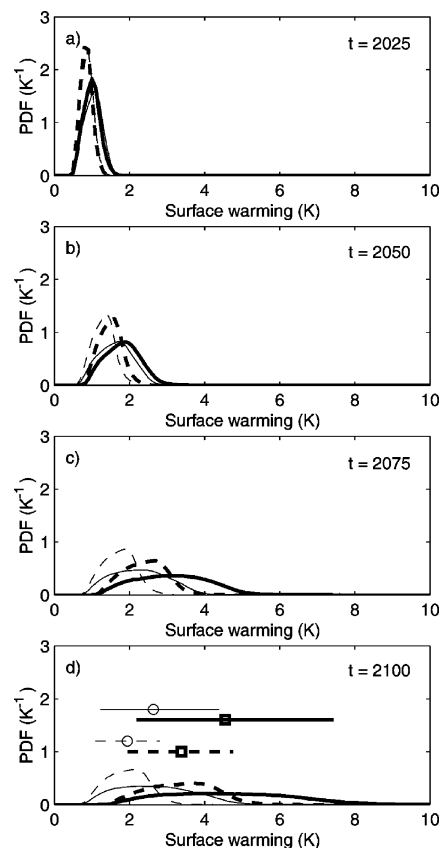


Fig. 6a–d PDFs for surface warming (relative to the average 1960–1990) for scenarios B1 (thin lines) and A2 (thick lines) for climate sensitivity constrained by observations (solid) and for IPCC climate sensitivities 1.5–4.5 K (dashed). Time slices are year 2025, 2050, 2075, 2100 in a to d. The horizontal lines in panel d indicate the medians and 5% to 95% probability ranges for the corresponding PDFs

PDFs get wider towards the end of the century. Although the contribution to the radiative forcing of the uncertain aerosol forcing decreases with time and is dominated by relatively well-known greenhouse gas components, the net uncertainty still grows monotonically with time due to uncertainties in climate sensitivity and the carbon cycle. If all climate sensitivities are considered (solid lines) as constrained in Fig. 3, the best-guess (median) values obtained for the year 2100 are 2.6 K and 4.6 K for B1 and A2, respectively. This is up to 30% higher than the numbers obtained by IPCC (2001), who agreed on best-guess values of 2.0 K and 3.8 K, based on several comprehensive AOGCMs. Further, we find probabilities of 51% for B1 and 45% for A2 that the warming exceeds the upper limit of the IPCC uncertainty range. The probability for a warming below the IPCC range, however, is small (around 10%). In absolute numbers, the discrepancy is, of course, particularly strong for the range of scenario A2, where we find an upper limit of 7.4 K for the 5% to 95% range, compared to 4.8 K found by IPCC. The reason for these large differences lies in the assumptions about climate sensitivity. If we adopt the range of 1.5 K to 4.5 K assumed by IPCC, our ranges (dashed lines) are broadly consistent with the ranges given by IPCC. We can therefore interpret the uncertainties proposed by IPCC as approximately covering the 5% to 95% probability range, if an upper limit of 4.5 K for climate sensitivity is assumed. However, higher climate sensitivities cannot be excluded by the observed atmospheric and oceanic warming. If higher climate sensitivities are considered, IPCC strongly underestimates the possibility for a warming largely exceeding the ranges published in the TAR.

In a recent paper, Stott and Kettleborough (2002) presented model simulations indicating that for the next few decades, the best-guess as well as the uncertainties for the projected warming depend only weakly on the choice of the scenario, both because of the delayed response of the Earth’s climate system to external forcing and because the differences between the individual scenarios in the first few decades are relatively small. The results shown in Fig. 6 confirm these findings. The differences between the scenarios are negligible for the first half of the next century. However, this is no longer the case after 2050, when the greenhouse gas emissions start to decrease for scenario B1, while they increase continuously until the end of the century for A2.

3.4 Sensitivity to consistency criteria

Four criteria are used here to test the consistency of model simulations with data. First, the modelled surface warming trend over the last century is required to match the observational trend within one standard deviation of the data. Similarly, the trend in the ocean heat uptake between 1955 and 1995 is required to match the observational trend. To guarantee a certain degree of

correlation over time, the absolute difference of the simulated and the observed surface warming is divided by the uncertainty of the observations at the corresponding time, and this quantity is then averaged over time. A maximum value of 1 for surface warming and 2 for ocean heat uptake is set for this dimensionless quantity, expressing the time-dependent agreement of model and data in terms of data uncertainties. In other words, the mean deviation of model and observations must not exceed one standard deviation of the data for the surface warming. For the ocean heat uptake, the agreement of model and data is significantly worse (see Subsect. 2.5), and agreement is required within two standard deviations. The choice of the criteria is not critical for the results. To test this, we calculated the surface warming at year 2020 for scenario B1 using only one of the four or different combinations of two consistency criteria. The atmospheric trend is found to be the strongest individual constraint, because its uncertainty is small. Both the oceanic trend and time-dependent correlation will constrain the ensemble significantly less, because of the large uncertainty and the short time period where observations are available. In general, the sensitivity of the results on the consistency criteria used is rather small. In general, about half of the simulations are consistent with one criterion, and only about 10% are consistent with all four criteria.

The inclusion of warming patterns instead of global averages could possibly improve the ensemble method. However, this would require models with higher resolution, and natural variability would have to be taken into account, since many patterns detected in the trend also occur due to internal variability. The inclusion of further datasets, e.g. upper air temperature, probably has little potential to improve the results (Forest et al. 2002). However, a longer record of ocean temperature with smaller uncertainty would help further constrain the model responses as well as ocean model versions.

An interesting point is that the projected warming as constrained either by the surface warming, or by the ocean heat uptake, are similar. The two observational datasets are consistent on time scales of more than a few decades in the sense that they suggest a comparable future warming.

3.5 Sensitivity to ocean model parameters

To account for the uncertainties in the ocean subgrid-scale mixing processes dominating the oceanic heat uptake, ten different ocean model parameterizations have been used here. We calculated the percentage of simulations consistent with data to test whether some of them are more realistic than others, and found the differences to be negligible. Using global values of observations, no ocean model version can be excluded, although the parameter ranges covered by the different versions are large. Further, to check the sensitivity of results to the different versions used, the surface warming and its uncertainty for the year 2100 and scenario B1

were calculated individually for each model version. In general, the differences are small, and difficult to attribute to the parameters used.

3.6 Quantifying the carbon cycle-climate feedback

The SRES illustrative scenarios provided by IPCC (2001) prescribe the emissions of greenhouse gases and aerosols for the twenty-first century. A model of the global carbon cycle is therefore applied first to calculate atmospheric greenhouse gas concentrations from emissions. In our case, this is done externally by calculating atmospheric concentrations under the assumption of constant climate. A variable feedback factor in the total radiative forcing is then used to account for the impact of warming on the carbon cycle (see Sect. 2.1 for details). This model setup allows us to quantify the “carbon cycle-climate feedback”, i.e. the difference of the projections obtained when assuming a constant climate for the carbon cycle compared to the case where the warming feedback on the carbon cycle is included. The resulting PDFs for the projected warming are shown in Fig. 7 for

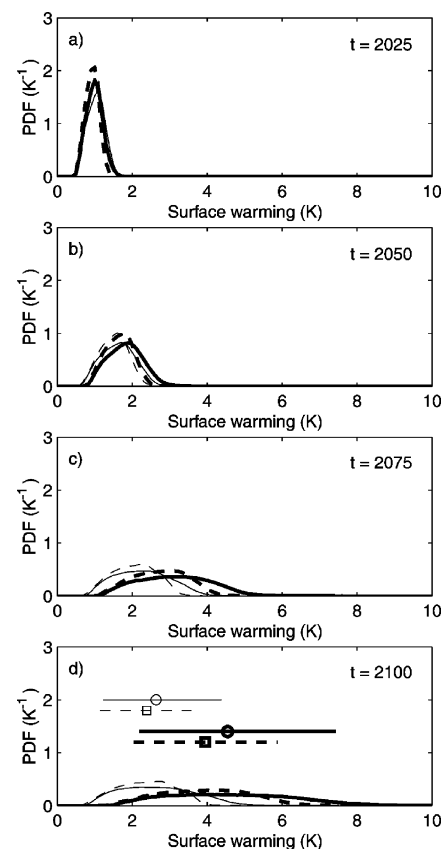


Fig. 7a–d PDFs for surface warming at year 2100 (relative to the average 1960–1990) for scenarios B1 (*thin lines*) and A2 (*thick lines*), for a carbon cycle under constant climate (*dashed*) and when including climate-carbon cycle feedback (*solid*). Climate sensitivity is only constrained by observations. Time slices are for the year 2025, 2050, 2075, 2100 in **a** to **d**. The *horizontal lines* in **d** indicate the medians and 5% to 95% probability ranges for the corresponding PDFs

the scenarios B1 (thin lines) and A2 (thick lines). The shift towards higher temperatures (dashed lines to solid lines) indicates the strength of the “carbon cycle-climate feedback”. While the differences are obviously small for the early decades and for low emission scenarios, they amount to considerable increases of 0.6 K in the A2 best-guess and 1.6 K for an upper limit of the A2 uncertainty range. Earlier studies have shown that the effect of a reduction in oceanic carbon uptake due to global warming on atmospheric CO_2 is relatively small, as long as no major reorganizations in the circulation occur (Joos et al. 1999; Plattner et al. 2001). The reduction in terrestrial carbon uptake dominates the changes in atmospheric CO_2 due to global warming and is mainly due to increased soil respiration and forest dieback (Meyer et al. 1999; Cox et al. 2000; Friedlingstein et al. 2000; Joos et al. 2001). The uncertainty of the carbon cycle is taken into account here by choosing different values for the carbon cycle feedback parameter γ in each simulation. However, when neglecting this uncertainty and keeping γ constant in all simulations, we find very similar results. Thus, the contribution of the carbon cycle feedback uncertainty to the total uncertainty in the projections is small.

3.7 Changes in the thermohaline circulation

The future warming of the climate system has the potential to weaken the Atlantic thermohaline circulation (THC) by reducing surface water density in the formation regions of North Atlantic Deep Water (NADW) through high-latitude warming and enhanced poleward moisture transport of the atmosphere (Manabe and Stouffer 1993; IPCC 2001). An important related issue is the question whether the climate system could cross a critical threshold leading to a non-linear transition of the THC to a qualitatively different circulation mode without deepwater formation in the North Atlantic. Such a reorganization would have a profound impact on the climate of at least the Atlantic region, on ocean biogeochemical cycles (Joos et al. 1999) and future sea level (Knutti and Stocker 2000). Most comprehensive ocean–atmosphere models show a significant reduction of NADW formation (see Rahmstorf 1999; IPCC 2001, and references therein) for a warming climate. For a strong warming, some models even simulate a complete shutdown of the THC (Manabe and Stouffer 1993; Stocker and Schmittner 1997). However, recent studies also suggest that the THC might be stabilized through increased wind-driven evaporation (Gent 2001) and trends in the North Atlantic Oscillation (Delworth and Dixon 2000) or the El Niño–Southern Oscillation (Latif et al. 2000; Schmittner et al. 2000). The stability of the THC in climate models depends on many factors (Stocker et al. 2001), among them the response of the hydrological cycle to the warming, the parameterization of mixing processes in the ocean (Knutti et al. 2000), the rate of greenhouse gas increase (Stocker and Schmittner

1997) and the initial strength of the THC itself (Tziperman 2000).

The Third Assessment Report of IPCC concludes that a reduction of the Atlantic THC is a likely response to increased greenhouse gas forcing, but that during this century, a complete shutdown is less likely than previously assumed. However, if the rate of change in radiative forcing remains large and lasts long enough, a complete shutdown after the year 2100 cannot be excluded. Unfortunately, the limited number of simulations available for the IPCC TAR and their disagreement again prevents any quantitative estimate of probabilities. To fill this gap, we have analyzed the response of the THC to global warming in our ensemble simulations and calculated a PDF for the overturning reduction projected for the next century. We find a very broad range of possible responses due to the uncertainties considered and due to different dynamical responses of the various ocean model versions. The median and 5% to 95% probability range of the relative THC strength for all model versions is shown in Fig. 8. Every model version has a different steady-state overturning rate, hence all results are scaled to their preindustrial overturning rate. While the changes and uncertainties over the industrial period are on the order of only 10%, the reduction in the year 2100 might be anywhere between 15% and 85% compared to the equilibrium state, with a best-guess reduction of about 65%. One might question the accuracy of the dynamical ocean circulation response in a coarse-resolution zonally averaged ocean model. However, these results are entirely in line with most of the comprehensive AOGCMs that show a reduction of the THC during the twenty-first century, but disagree in how strong the reduction should be.

3.8 Steric sea level rise

There is evidence from observations and models that sea level has risen during the last one hundred years by about 1 mm y^{-1} due to increasing surface air and ocean

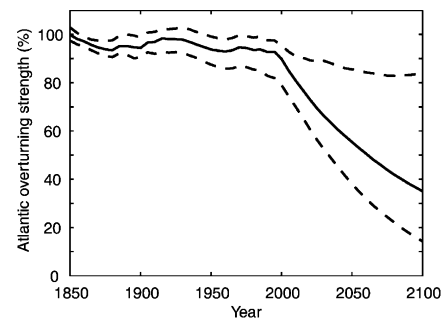


Fig. 8 Atlantic overturning strength relative to the preindustrial steady state for scenario B1, scaled to preindustrial equilibrium. All ten model versions are included and fitted with neural networks. The *solid line* denotes the median, the *dashed lines* indicate the 5 to 95% probability range. The results for scenario A2 are very similar and have been omitted for clarity

temperatures (see Church and Gregory 2001, for a review). Future warming is expected to cause an acceleration in sea level rise until the year 2100 and beyond. Even if the emissions of greenhouse gases were stabilized and reduced, sea level would continue to rise for hundreds of years due to the slow mixing of heat into the ocean. The main anthropogenic factors contributing to changes in sea level are changes related to oceanic thermal expansion, glaciers and small ice caps, the large ice sheets of Greenland and Antarctica and possible changes in ground-water storage. The IPCC (2001) has agreed on a projected range for sea level rise at year 2100 of 0.09 to 0.88 m for all SRES scenarios, with a central value of 0.48 m. About half of the anthropogenically caused sea level rise can be attributed to the thermal expansion of the seawater.

Figure 9 shows the 5% to 95% probability range for steric sea level rise as simulated by our ensemble method for the scenarios B1 (thin lines) and A2 (thick lines). The best guess values for the year 2100 are estimated at 0.48 m for B1 and 0.62 m for A2, with uncertainties of roughly ± 0.2 m in both cases. The median values are significantly larger than those of IPCC, again due to the inclusion of high climate sensitivities.

4 Discussion and conclusions

While climate models have been improved significantly in recent years with regard to resolution and simulated processes, the question of the uncertainty in model simulations has often been treated unsatisfactorily. We believe that this is at least partly due to the lack of a consensus about how uncertainties should be quantified, and partly because statistical estimates of uncertainty require a considerable number of simulations that have usually not been available. The idea of scaling future projections by comparing fingerprints of the simulated and observed warming patterns from recent decades (Allen et al. 2000, 2002) provided the first objective statistical uncertainty estimates for future warming. The fingerprint scaling method accurately

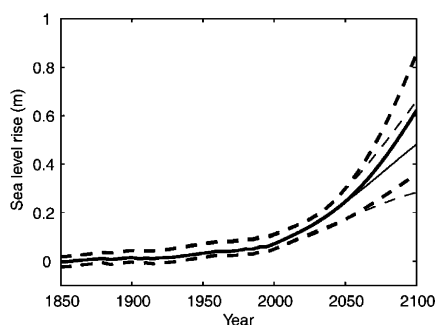


Fig. 9 Projected changes in sea level for the scenarios B1 (thin lines) and A2 (thick lines), using all model versions. The solid lines denote the median, the dashed lines indicate the 5% to 95% probability ranges of the ensemble

takes into account internal variability and does not rely on an estimate for climate sensitivity. However, not all uncertainties (e.g. in the model configuration) can be included systematically, and projections are only valid as long as the ratio of greenhouse gas warming and aerosol cooling remains similar to the current value, because the two patterns cannot be detected separately with sufficient accuracy.

The results presented in this study use a different approach and constrain the uncertainties from a large number of ensemble simulations that are consistent with the observed surface warming and ocean heat uptake of the last decades and centuries. The restrictions outlined do not apply to this method, and in principle, any kind of uncertainty in model parameters, model configuration and resolution, parameterizations, observational data-sets etc., can be taken into account. For climate projections covering the next 100 years, the most important are the uncertainties in climate sensitivity, in the radiative forcing, in the carbon cycle, and in oceanic heat uptake. Considering all of these main uncertainties, we find a probability of about 40% that the surface warming at year 2100 exceeds the uncertainty range proposed by IPCC (2001), but only one of about 10% for the warming to be lower than the range of IPCC (2001). While the best-guess values are only 20% to 30% higher, the upper limits of the uncertainty ranges exceed those of IPCC by 50% to 70%, depending on the scenario. The main reason is that the climate sensitivity is only poorly constrained by the observational data. In particular, high values for the climate sensitivity combined with a strong negative aerosol radiative forcing can reasonably reproduce the observed warming and can therefore not be excluded. The choice of the a priori upper limit for climate sensitivity has a strong influence on the upper limit of the warming range determined by this method. Climate sensitivities of more than 10 K are difficult to produce in our model using the feedback parameterization described in Sect. 2.1 (possibility of runaway greenhouse effects for high values of λ). A value of 10 K is chosen here to allow for high climate sensitivities, but avoid unreasonable results by the climate model. If the range of climate sensitivities of 1.5 K to 4.5 K adopted by IPCC (2001) is considered, the surface warming uncertainties published in the IPCC TAR approximately cover the 5% to 95% confidence ranges derived with the ensemble method.

Based on the parameterization introduced here to account for the feedback of reduced oceanic and terrestrial carbon uptake and global warming, we estimate an additional increase in the best guess surface warming for year 2100 of 10% to 15% (20% to 30% for the upper limit of the uncertainty range) compared to the case when CO_2 concentrations are calculated from emissions under constant climate. While this feedback between global warming and reduced carbon uptake is significant, the uncertainty of the feedback is relatively small compared to the uncertainties in radiative forcing and in climate sensitivity.

In addition to surface warming, sea level from thermal expansion and the change in the thermohaline circulation (THC) can be projected for the next century. The ensemble model projects a THC reduction of 15% to 85% at year 2100, relative to a preindustrial equilibrium state. This is a broad range for one scenario, especially since all model versions have similar dynamics, resolution, and the same atmospheric model component. But it is in agreement with the large spread of projected THC responses summarized in the IPCC TAR.

The ensemble method can also be used to constrain parameters in the climate system. Whereas the climate sensitivity and the rate of ocean heat uptake are only poorly constrained by observations, the uncertain indirect aerosol forcing component (plus any forcing not explicitly considered) can be narrowed considerably. A range of 0 to -1.2 W m^{-2} is found to be consistent with the observed warming of ocean and atmosphere.

Unfortunately the method used here does not consider internal variability in the climate system. Thus our results rely on the assumption that external solar, volcanic and anthropogenic radiative forcing is sufficient to explain the observed warming on decadal scales. Internal variability in the climate system may mask the externally forced changes, and external forcings are even anticipated to induce trends or regime shifts in natural patterns of preferred weather situations (Corti et al. 1999), in the El Niño-Southern Oscillation (Timmermann et al. 1999; Latif et al. 2000), or the North Atlantic Oscillation (Delworth and Dixon 2000), making it hard to distinguish externally forced changes from natural variability. However, the Third Assessment Report of IPCC (2001) concludes, based on detection and attribution studies, that there is strong evidence that most of the observed warming is caused by anthropogenic greenhouse gas and aerosol emissions. An independent quantitative estimate from three 1000 yr control simulations of different AOGCMs (Stouffer et al. 2000) yields a probability of 5% for a 100 yr trend in global mean surface temperature larger than about 0.11 K, and a probability of 1% for a 100 yr trend larger than about 0.15 K produced by unforced internal variability alone. If these models are roughly correct in simulating the magnitude of natural variability, we could expect a possible unforced trend over the last century due to internal variability to be smaller than the uncertainties in the observed temperature trend over the last century, which is estimated at about 0.2 K (IPCC 2001). Further confidence in the results is given by the fact that two independent studies, one using a fingerprint scaling method with a comprehensive AOGCM (Stott and Kettleborough 2002), the other using an ensemble method with a simplified climate model (Knutti et al. 2002), estimated almost identical best-guess values and similar uncertainty ranges for the next few decades and an illustrative SRES scenario.

The results presented here are only weakly sensitive to the a priori assumptions for the forcing uncertainties and to parameter settings affecting the ocean heat

uptake. The projected warming, however, is sensitive to the assumed range for climate sensitivity, as discussed.

The available computational resources currently prevent large ensemble simulations with complex models. As long as simple models are used in the ensemble method, which do not simulate internal variability, tools for function approximation can be used to improve the efficiency of the ensemble method. The approach presented here uses a neural network to directly predict relevant model output from an input parameter set, without explicitly integrating the climate model forward in time, and has shown to be fast and accurate. The ensemble size can thus easily be increased to 10^6 or more, and extensive sensitivity studies are feasible.

The approach presented here can be seen as a synthesis trying to bring several independently estimated but strongly linked quantities together within the framework of one climate model. On the long-term, future progress can only be achieved by a better understanding of the processes in the climate system. Apart from progress in the ability of climate models to correctly represent the relevant processes in the climate system, three main points can be identified which would reduce the uncertainties of future ensemble climate projections, better constrain parameters and processes in climate models. Longer and more accurate reconstructions of the recent warming, in particular in the ocean, would tighten the constraints on the ensemble simulations. Further, reducing the uncertainties in either climate sensitivity or the aerosol radiative forcing, or both, would strongly reduce the uncertainties in the projections. Finally, increased computational power would allow large ensembles using more comprehensive climate models, which simulate natural variability and consider more feedback processes in greater detail. In addition, regional patterns of the observed climate change could be used more effectively with high-resolution models. The fact that each member of an ensemble is independent of the other simulations, makes this method well suited for relatively inexpensive distributed computing environments. The ensemble method is therefore useful to objectively estimate uncertainties in the climate system and in projections of future climate change.

Acknowledgements We thank T. Crowley for providing the volcanic and solar radiative forcing reconstructions and S. Levitus for compiling the time series of ocean heat uptake. We enjoyed discussions with S. Gerber, N. Edwards and J. Flückiger. This work was supported by the Swiss National Foundation.

Appendix 1: estimating the carbon cycle-climate feedback factor

To account for the uncertainty in the carbon cycle in the individual model simulations, a feedback factor γ was introduced, describing the feedback of a warming climate on the oceanic and terrestrial carbon cycle (see Sect. 2.2).

The value of the feedback factor γ and the baseline for CO_2 radiative forcing were derived separately from simulations with

the Bern Carbon Cycle-Climate (Bern CC) model (Joos et al. 2001; Gerber et al. 2003). Details of the model and of the simulations are described elsewhere (Joos et al. 2001). In the Bern CC model, the physical climate system is represented by an impulse response-empirical orthogonal function substitute of the ECHAM3/LSG atmosphere/ocean general circulation model (Hooss et al. 2001) which is driven by radiative forcing. The impulse response function and the spatial patterns (EOFs) of the simulated changes in temperature, precipitation and cloud cover were derived from an 800 year long ECHAM3/LSG AOGCM simulation wherein atmospheric CO_2 was quadrupled in the first 140 years and held constant thereafter (Voss and Mikolajewicz 2001). The climate sensitivity of the substitute is prescribed. The carbon cycle module includes a well-mixed atmosphere, the HILDA ocean model (Siegenthaler and Joos 1992; Joos et al. 1996), and the Lund-Potsdam-Jena Dynamic Global Vegetation Model (LPJ-DGVM) (Sitch et al., Submitted 2002). The effect of sea surface warming on carbon chemistry is included (Joos et al. 1999), but ocean circulation changes are not considered. The LPJ-DGVM is run on a resolution of $3.75^\circ \times 2.5^\circ$ and is driven by local temperatures, precipitation, incoming solar radiation, cloud cover, and atmospheric CO_2 . It simulates the distribution of nine plant functional types based on bioclimatic limits for plant growth and regeneration. Photosynthesis is a function of absorbed photosynthetically active radiation, temperature, atmospheric CO_2 concentration and an empirical convective boundary layer parameterization to couple the carbon and water cycles. After spinup, the model was driven by prescribed anthropogenic radiative forcing and observed atmospheric CO_2 for the period from 1765 to 2000. Afterwards, the atmospheric concentration of CO_2 and other agents and radiative forcing were calculated from emissions of greenhouse gases and aerosol precursors. The baseline trajectory for CO_2 forcing was calculated from simulations where the climate sensitivity was set to zero. The feedback parameter γ is calculated from the difference in CO_2 radiative forcing between simulations with and without global warming divided by the simulated temperature increase since year 2000 of the warming simulation. It varies around $0.25 \text{ W m}^{-2}/\text{K}$ for a range of emission scenarios and for a range of climate sensitivities (see Fig. 10). It is noted that the feedback factor γ as derived here implicitly accounts for the climate change up to year 2000; γ would be somewhat lower when derived from simulations where emissions are prescribed over the historical period and expressed relative to the warming since preindustrial time. The uncertainty of γ has been estimated based on the simulations of different scenarios with the model described and based on other published modelling studies (Friedlingstein et al. 2000; Cox et al. 2000). The feedback strength in the IPSL model (Friedlingstein et al. 2000) is slightly lower than in the Bern CC model, whereas the Hadley centre model has a much stronger feedback (Cox et al. 2000).

Appendix 2: neural network design

1.1 Principles of neural networks

The research field of neural networks is relatively young and has experienced three periods of extensive activity. The first wave of interest emerged after the pioneering work of McCulloch and Pitts in 1943. The second occurred in the 1960s with the *Perceptron Convergence Theorem* of Rosenblatt (1962), and was damped by the work of Minsky and Papert (1969) showing the limitations of a simple perceptron. In the early eighties, new theoretical results (e.g. the Hopfield energy approach and the discovery of error back-propagation), and increased computational capacities renewed the interest in neural networks. They are now used to solve a variety of problems in pattern recognition and classification, tracking, control systems, fault detection, data compression, feature extraction, signal processing, optimization problems, associative memory and more.

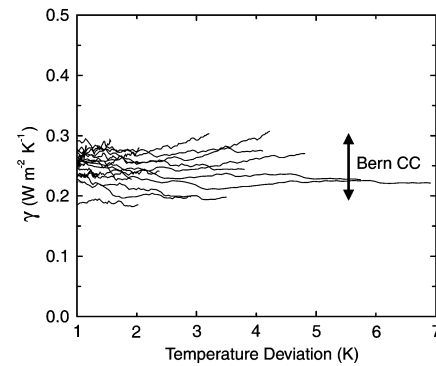


Fig. 10 Global warming-carbon cycle feedback factor γ versus global mean surface warming since year 2000 AD. The feedback factors have been evaluated for the six illustrative SRES scenarios and applying climate sensitivities of 1.5, 2.5, and 4.5 K in the Bern Carbon Cycle-Climate Model (Joos et al. 2001). Atmospheric CO_2 , its forcing and radiative forcing by other agents were prescribed up to year 2000 according to observations. Afterwards, concentrations and radiative forcing were simulated from emissions of greenhouse gases and aerosol precursors. Simulations with global warming and baseline simulations without global warming (climate sensitivity set to 0 K) were performed. The feedback factors were computed from the model output for the period 2000 to 2100. First, the difference in CO_2 radiative forcing between a global warming simulation and its baseline simulation was determined; then this difference was divided by the increase in global mean surface temperature realized since year 2000 in the warming simulation. Only results for temperature changes larger than 1 K are shown, i.e. when the factors have converged to a relatively stable value

There are many different types of neural networks, each suitable for specific applications. The L -layer feed-forward network consists of one input layer, $(L - 2)$ hidden layers, and one output layer of units successively connected in a feed-forward fashion with no connections between units in the same layer and no feedback connections between layers. A sufficiently large three-layer feed-forward network can approximate any functional relationship between inputs and outputs. We use such a three-layer feed-forward network, with $N = 10$ units in the input and hidden layer. The number of neurons in the output layer is equal to the number of output values desired. The choice of the network size N is critical, and discussed in Appendix 1.2 for our application.

Before the neural network can perform anything reasonable, it has to go through a training or learning phase, where the connection weights w_j (and possibly also the network architecture) are continuously updated so that the network can efficiently perform a specific task. The ability of neural networks to learn automatically from examples makes them attractive and exiting. Instead of following a set of rules specified by human experts, the neural network appears to learn the underlying input-output relationship from the examples presented to it in a training set. There are numerous learning algorithms, which can be classified into supervised and unsupervised learning. In our case, a limited training set of parameter vectors \vec{R} and the corresponding simulations $C(\vec{R}, t)$ from the climate model are available. In this case, the supervised learning phase is the problem of minimizing the sum of the squared differences between the climate model simulations $C(\vec{R}, t)$ and the corresponding substitute response $F(\vec{R}, t)$ approximated by the neural network by adjusting the weights w_j in each of the neurons. For efficiency, only the desired subset of the model output $C(\vec{R}, t)$ is approximated. The Levenberg-Marquardt (Hagan and Menhaj 1994) algorithm is the most efficient learning algorithm for our application due to its extremely fast convergence. A detailed introduction to neural network architectures, learning rules, training methods and applications was written by Hagan et al. (1996).

Fig. 11a–d Mean error and correlation coefficient between climate model output and neural network predictions for surface warming and ocean heat uptake as a function of the network size. The network consists of three layers, and there are N neurons in layer 1 and 2. The number of neurons in the output layer 3 is fixed and given by the output size

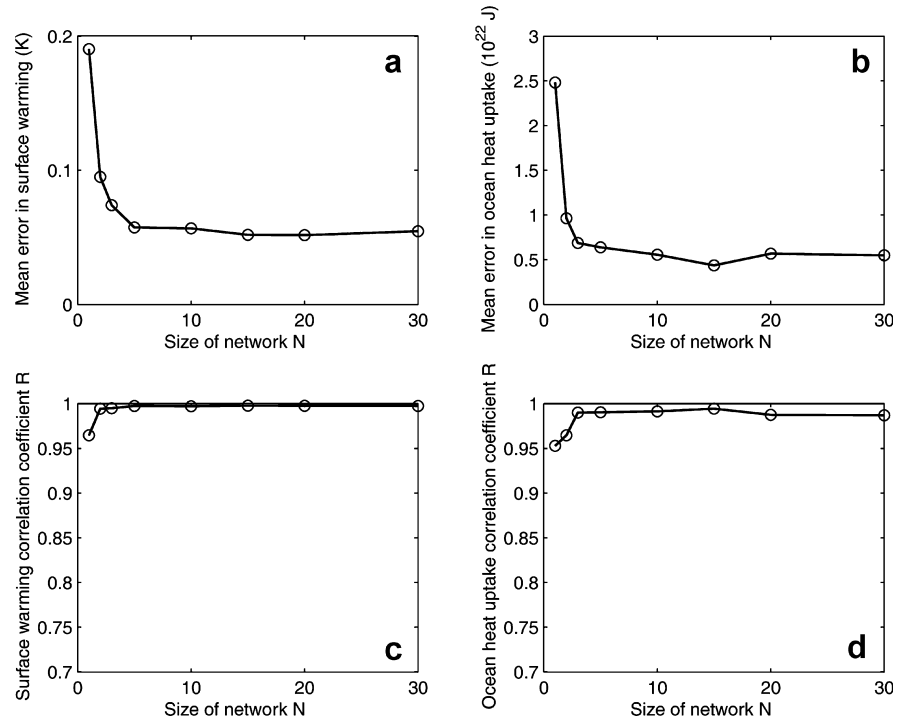
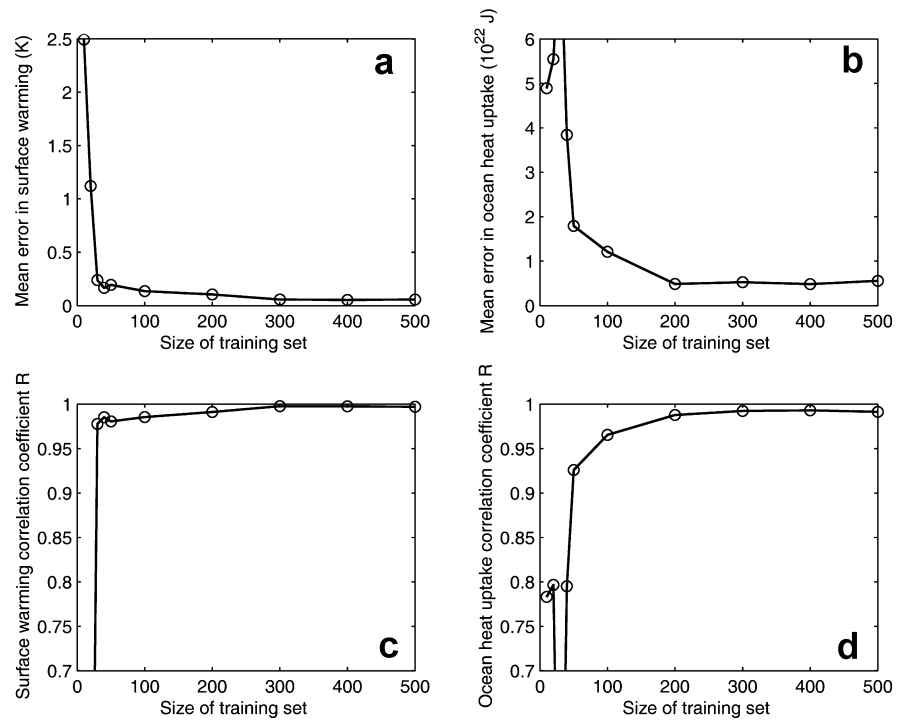


Fig. 12a–d Mean error and correlation coefficient between climate model output and neural network predictions for surface warming and ocean heat uptake as a function of the size of the training set



1.2 Sensitivity to neural network size

The size of the neural network, i.e. the number of neurons in the input and hidden layer, is critical for performance as well as for efficiency reasons. If the network is too small, it has too few degrees of freedom to learn the desired behaviour and will do a poor approximation of the input–output relationship. If it is too large, it will be computationally expensive to run. Further, too many neurons can contribute to “overfitting”, the case in which all training samples are very well fit, but the fitting curve takes wild oscillations

between these points, causing the network to lose its ability to generalize and correctly predict the output from an input it has not been trained with. One way to accomplish this task is to start with a small network and to subsequently increase the number of neurons in the layers. The mean error will drop asymptotically to zero or to an approximately constant value if the function cannot be approximated perfectly (which is the case here). This is shown in Fig. 11a, b for the mean error (difference between climate model and neural network) of the surface warming and ocean heat uptake over the observational period. Further, correlation coefficients

between the approximated function and the climate model-predicted time evolution are shown (Fig. 11c, d). The choice of $N = 10$ neurons in the input and hidden layer is reasonable, since a larger network does not perform better, but takes significantly longer to train. The number of neurons in the input and hidden layers must not necessarily be identical. Sigmoid and linear activation functions were used in the input and hidden layer, respectively. Other activation functions yield similar or less accurate results (not shown).

1.3 Sensitivity to size of training set

Choosing an appropriate set of training patterns for the neural network is the second important point. Too few or poorly selected training samples will prevent the network from learning the whole input–output relationship, and only part of the presented input will be recognized correctly. Too many training samples, on the other hand, will not do any harm, but increase the computational cost dramatically. Again, the performance of the network was evaluated here for different sizes of the training sample set (Fig. 12). A training set of 500 simulations was found to be still efficient enough to handle and was chosen to ensure proper training over the whole parameter space. An interesting feature is seen in Fig. 12b, d for small numbers of training samples. The mean error does not necessarily decrease when adding more training samples, if the samples are unevenly distributed in parameter space. Adding a few samples lying close together in parameter space by coincidence causes the network to behave better in this specific region, but poorly at other locations. If the samples are randomly chosen, then only a sufficiently large sample will ensure an approximately even distribution and good performance in the whole parameter range of interest.

The problem of overfitting mentioned above can be tackled in several ways. Here we adopt a method called “early stopping”, which ensures the generalization ability by stopping the training process before overfitting starts. This is achieved by using a separate validation set instead of the training set to test the performance of the network. Typically, the validation error will decrease during the initial phase of the training, as does the training set error. However, when the network begins to overfit the data, the error on the validation set will usually begin to rise, while the error on the training set is still decreasing. When the validation error increases for a specified number of iterations, the training is stopped, and the network weights at the minimum of the validation error are returned.

References

- Allen MR, Gillett NP, Kettleborough JA, Hegerl G, Schnur R, Stott PA, Boer G, Covey C, Delworth TL, Jones GS, Mitchell JFB, Barnett TP (2002) Quantifying anthropogenic influence on recent near-surface temperature change. *Surv Geophys* (In press)
- Allen MR, Stott PA, Mitchell JFB, Schnur R, Delworth TL (2000) Quantifying the uncertainty in forecasts of anthropogenic climate change. *Nature* 407: 617–620
- Andronova N, Schlesinger ME (2001) Objective estimation of the probability distribution for climate sensitivity. *J Geophys Res* 106: 22,605–22,612
- Barnett TR, Pierce DW, Schnur R (2001) Detection of anthropogenic climate changes in the World’s oceans. *Science* 292: 270–274
- Boucher O, Haywood J (2001) On summing the components of radiative forcing of climate change. *Clim Dyn* 18: 297–302
- Church JA, Gregory JM (2001) Changes in sea level. *Climate change 2001: the scientific basis*. In: Houghton JT et al. (eds) *Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK, pp 639–693
- Corti S, Molteni F, Palmer TN (1999) Signature of recent climate change in frequency of natural atmospheric circulation regimes. *Nature* 398: 799–802
- Cox P, Betts R, Jones C, Spall S, Totterdell I (2000) Acceleration of global warming due to carbon-cycle feedbacks in a coupled climate model. *Nature* 408: 184–187
- Crowley TJ (2000) Causes of climate change over the past 1000 years. *Science* 289: 270–277
- Delworth TL, Dixon KW (2000) Implications of the recent trend in the Arctic/North Atlantic Oscillation for the North Atlantic thermohaline circulation. *J Clim* 13: 3721–3727
- Forest CE, Stone PH, Sokolov AP, Allen MR, Webster MD (2002) Quantifying uncertainties in climate system properties with the use of recent climate observations. *Science* 295: 113–117
- Friedlingstein P, Bopp L, Ciais P, Fairhead JLDL, LeTreut H, Monfray P, Orr J (2000) Positive feedback of the carbon cycle on future climate change. *Tech Rep 19, Note du Pole de Modelisation, Institute Pierre Simon Laplace, France*
- Gent PR (2001) Will the North Atlantic Ocean thermohaline circulation weaken during the 21st century? *Geophys Res Lett* 28: 1023–1026
- Gent PR, McWilliams JC (1990) Isopycnal mixing in ocean circulation models. *J Phys Oceanogr* 20: 150–155
- Gent PR, Willebrand J, McDougall TJ, McWilliams JC (1995) Parameterizing eddy-induced tracer transports in ocean circulation models. *J Phys Oceanogr* 25: 463–474
- Gerber S, Joos F, Brügger P, Stocker TF, Mann ME, Sitch S, Scholze M (2003) Constraining temperature variations over the last millennium by comparing simulated and observed atmospheric CO₂. *Clim Dyn* 20: 281–299, doi: 10.1007/s00382-002-0270-8
- Gregory JM, Stouffer RJ, Raper SCB, Stott PA, Rayner NA (2002) An observationally based estimate of the climate sensitivity. *J Clim* 15: 3117–3121
- Hagan MT, Menhaj M (1994) Training feedforward networks with the Marquart algorithm. *IEEE Trans Neural Networks* 5: 989–993
- Hagan MT, Demuth HB, Beale MH (1996) *Neural Network Design*. PWS Publishing, Boston, USA
- Hansen JE, Sato M, Lacis A, Ruedy R, Tegen I, Matthews E (1998) Climate forcings in the Industrial era. *Proc US Natl Acad Sci* 95: 12,753–12,758
- Hegerl GC, Stott PA, Allen MR, Mitchell JFB, Tett SFB, Cubasch U (2000) Optimal detection and attribution of climate change: sensitivity of results to climate model differences. *Clim Dyn* 16: 737–754
- Hooss G, Voss R, Hasselmann K, Maier-Reimer E, Joos F (2001) A nonlinear impulse response model of the coupled carbon cycle-climate system (NICCS). *Clim Dyn* 18: 189–202
- IPCC (2001) *Climate change 2001: the scientific basis*. In: Houghton JT et al. (eds) *Contribution of Working Group I to the Third Assessment Report of the Intergovernmental Panel on Climate Change*. Cambridge University Press, Cambridge, UK, pp 881
- Jones PD, New M, Parker DE, Martin S, Rigor I (1999) Surface air temperature and its changes over the past 150 years. *Rev Geophys* 37: 173–199
- Joos F, Bruno M (1996) Pulse response functions are cost-efficient tools to model the link between carbon emissions, atmospheric CO₂ and global warming. *Phys Chem Earth* 21: 471–476
- Joos F, Bruno M, Fink R, Stocker TF, Siegenthaler U, Le Quéré C, Sarmiento JL (1996) An efficient and accurate representation of complex oceanic and biospheric models of anthropogenic carbon uptake. *Tellus Ser B* 48: 397–417
- Joos F, Plattner GK, Stocker TF, Marchal O, Schmittner A (1999) Global warming and marine carbon cycle feedbacks on future atmospheric CO₂. *Science* 284: 464–467
- Joos F, Prentice IC, Sitch S, Meyer R, Hooss G, Plattner GK, Gerber S, Hasselmann K (2001) Global warming feedbacks on terrestrial carbon uptake under the IPCC emission scenarios. *Global Biogeochem Cyc* 15: 891–907

- Knutti R, Stocker TF (2000) Influence of the thermohaline circulation on projected sea level rise. *J Clim* 13: 1997–2001
- Knutti R, Stocker TF (2002) Limited predictability of the future thermohaline circulation close to an instability threshold. *J Clim* 15: 179–186
- Knutti R, Stocker TF, Wright DG (2000) The effects of subgrid-scale parameterizations in a zonally averaged ocean model. *J Phys Oceanogr* 30: 2738–2752
- Knutti R, Stocker TF, Joos F, Plattner GK (2002) Constraints on radiative forcing and future climate change from observations and climate model ensembles. *Nature* 416: 719–723
- Latif M, Roeckner E, Mikolajewicz U, Voss R (2000) Tropical stabilization of the thermohaline circulation in a greenhouse warming simulation. *J Clim* 13: 1809–1813
- Leggett J, Pepper WJ, Swart RJ (1992) Climate change 1992. The Supplementary Report to the IPCC Scientific Assessment, Ch Emissions Scenarios for IPCC: an Update. Cambridge University Press, Cambridge, UK, pp 69–95
- Levitus S, Antonov JI, Boyer TP, Stephens C (2000) Warming of the World Ocean. *Science* 287: 2225–2229
- Levitus S, Antonov JI, Wang J, Delworth TL, Dixon KW, Broccoli AJ (2001) Anthropogenic warming of Earth's climate system. *Science* 292: 267–270
- Manabe S, Stouffer RJ (1993) Century-scale effects of increased atmospheric CO₂ on the ocean–atmosphere system. *Nature* 364: 215–218
- McCulloch WS, Pitts W (1943) A logical calculus of the ideas immanent in nervous activity. *Bull Math Biophys* 5: 115–133
- Meyer R, Joos F, Esser G, Heimann M, Kohlmaier GHG, Sauf W, Voss R, Wittenberg U (1999) The substitution of high-resolution terrestrial biosphere models and carbon sequestration in response to changing CO₂ and climate. *Global Biogeochem Cycles* 13: 785–802
- Minsky M, Papert S (1969) Perceptrons: an introduction to computational geometry. MIT Press, USA
- Myhre G, Highwood EJ, Shine KP, Stordal F (1998) New estimates of radiative forcing due to well mixed greenhouse gases. *Geophys Res Lett* 25: 2715–2718
- Nakićenović et al (2000) Special Report on Emission Scenarios. Intergovernmental Panel on Climate Change, Cambridge University Press, Cambridge, UK, pp 599
- Plattner GK, Joos F, Stocker TF, Marchal O (2001) Feedback mechanisms and sensitivities of ocean carbon uptake under global warming. *Tellus* 53B: 564–592
- Rahmstorf S (1999) Shifting seas in the greenhouse? *Nature* 399: 523–524
- Reichert BK, Schnur R, Bengtsson L (2002) Global ocean warming tied to anthropogenic forcing. *Geophys Res Lett* 29: 690–691
- Rosenblatt R (1962) Principles of neurodynamics. Spartan Books, New York, USA
- Santer BD, Taylor KE, Wigley TML, Johns TC, Jones PD, Karoly DJ, Mitchell JFB, Oort AH, Penner JE, Ramaswamy V, Schwarzkopf MD, Stouffer RJ, Tett S (1996) A search for human influences on the thermal structure of the atmosphere. *Nature* 382: 39–46
- Sarmiento JL, Le Quéré C (1996) Oceanic carbon dioxide uptake in a model of century-scale global warming. *Science* 274: 1346–1350
- Schmittner A, Stocker TF (1999) The stability of the thermohaline circulation in global warming experiments. *J Clim* 12: 1117–1133
- Schmittner A, Stocker TF (2001) A seasonally forced ocean model for paleoclimate studies. *J Clim* 14: 1055–1068
- Schmittner A, Appenzeller C, Stocker TF (2000) Enhanced Atlantic freshwater export during El Niño. *Geophys Res Lett* 27: 1163–1166
- Schneider SH (2001) What is 'dangerous' in climate change? *Nature* 411: 17–19
- Shine KP, Derwent RG, Wuebbles DJ, Morcrette JJ (1995) Radiative forcing of climate. In: Houghton JY, Jenkins GJ, Ephraums JJ (eds) Climate change, the IPCC scientific assessment, Cambridge University Press, Cambridge, UK
- Siegenthaler U, Joos F (1992) Use of a simple model for studying oceanic tracer distributions and the global carbon cycle. *Tellus Ser B* 44: 186–207
- Sitch S, Smith B, Prentice IC, Arneth A, Bondeau A, Cramer W, Kaplan J, Levis S, Lucht W, Sykes M, Thonicke K, Venevsky S (2002) Evaluation of ecosystem dynamics, plant geography and terrestrial carbon cycling in the LPJ Dynamic Global Vegetation Model. *Global Change Biol* (submitted)
- Stocker TF, Schmittner A (1997) Influence of CO₂ emission rates on the stability of the thermohaline circulation. *Nature* 388: 862–865
- Stocker TF, Wright DG (1991) A zonally averaged model for the thermohaline circulation. Part II: interocean exchanges in the Pacific-Atlantic basin system. *J Phys Oceanogr* 21: 1725–1739
- Stocker TF, Wright DG, Mysak LA (1992) A zonally averaged, coupled ocean–atmosphere model for paleoclimate studies. *J Clim* 5: 773–797
- Stocker TF, Knutti R, Plattner GK (2001) The future of the thermohaline circulation – a perspective. In: Seidov D, Haupt BJ, Maslin M (eds), The oceans and rapid climate change: past, present, and future, vol. 126 of Geophysical Monograph, American Geophysical Union, Washington D.C., USA
- Stott PA, Kettleborough JA (2002) Origins and estimates of uncertainty in predictions of twenty-first century temperature rise. *Nature* 416: 723–726
- Stott PA, Tett SFB, Jones GS, Allen MR, Mitchell JFB, Jenkins GJ (2000) External control of 20th century temperature by natural and anthropogenic forcing. *Science* 290: 2133–2137
- Stott PA, Tett SFB, Jones GS, Allen MR, Ingram WJ, Mitchell JFB (2001) Attribution of twentieth century temperature change to natural and anthropogenic causes. *Clim Dyn* 17: 1–21
- Stouffer RJ, Hegerl G, Tett S (2000) A comparison of surface air temperature variability in three 1000-yr coupled ocean–atmosphere model integrations. *J Clim* 13: 513–537
- Tett SFB, Stott PA, Allen MR, Ingram WJ, Mitchell JFB (1999) Causes of twentieth-century temperature change near the Earth's surface. *Nature* 399: 569–572
- Timmermann A, Oberhuber J, Bacher A, Esch M, Latif M, Roeckner E (1999) Increased El Niño frequency in a climate model forced by future greenhouse warming. *Nature* 398: 694–697
- Tziperman E (2000) Proximity of the present-day thermohaline circulation to an instability threshold. *J Phys Oceanogr* 30: 90–104
- Voss R, Mikolajewicz U (2001) Long-term climate changes due to increased CO₂ concentration in the coupled atmosphere–ocean general circulation model ECHAM3/LSG. *Clim Dyn* 17: 45–60
- Wigley TML, Raper SCB (2001) Interpretation of high projections for global-mean warming. *Science* 293: 451–454
- Wright DG, Stocker TF (1991) A zonally averaged ocean model for the thermohaline circulation, Part I: model development and flow dynamics. *J Phys Oceanogr* 21: 1713–1724
- Zwiers FA (2002) The 20-year forecast. *Nature* 416: 690–691