Clim. Past, 9, 1141–1152, 2013
www.clim-past.net/9/1141/2013/
doi:10.5194/cp-9-1141-2013
© Author(s) 2013. CC Attribution 3.0 License.

Climate of the Past
Open Access

An assessment of particle filtering methods and nudging for climate state reconstructions

S. Dubinkina and H. Goosse

Earth and Life Institute, Georges Lemaître Centre for Earth and Climate Research, Université catholique de Louvain, P.O. Box L4.03.07, 1348 Louvain-la-Neuve, Belgium

Correspondence to: S. Dubinkina (svetlana.dubinkina@uclouvain.be)

Received: 12 December 2012 – Published in Clim. Past Discuss.: 3 January 2013
Revised: 29 March 2013 – Accepted: 21 April 2013 – Published: 17 May 2013

Abstract. Using the climate model of intermediate complexity LOVECLIM in an idealised framework, we assess three data-assimilation methods for reconstructing the climate state. The methods are a nudging, a particle filter with sequential importance resampling, and a nudging proposal particle filter and the test case corresponds to the climate of the high latitudes of the Southern Hemisphere during the past 150 yr. The data-assimilation methods constrain the model by pseudo-observations of surface air temperature anomalies obtained from the same model, but different initial conditions. All three data-assimilation methods provide with good estimations of surface air temperature and of sea ice concentration, with the nudging proposal particle filter obtaining the highest correlations with the pseudo-observations. When reconstructing variables that are not directly linked to the pseudo-observations such as atmospheric circulation and sea surface salinity, the particle filters have equivalent performance and their correlations are smaller than for surface air temperature reconstructions but still satisfactory for many applications. The nudging, on the contrary, obtains sea surface salinity patterns that are opposite to the pseudo-observations, which is due to a spurious impact of the nudging on vertical exchanges in the ocean.

1 Introduction

Reliable reconstructions of the past climate states are essential for a comprehensive understanding of the climate system, more accurate climate predictions and projections. They enable to estimate the magnitude of the natural variability without anthropogenic impact and to provide insights into the processes responsible for climate changes.

A new but highly appealing approach to reconstruct the past climate states is data assimilation (e.g., Bhend et al., 2012; Widmann et al., 2010; Annan and Hargreaves, 2012). The main purpose of data assimilation is to estimate the state of a system as accurately as possible incorporating all the available information: numerical modelling of the behaviour of the system, observations and uncertainties of the model and of the observations (Talagrand, 1997). When choosing a data-assimilation method, the application to which it is applied should be kept in mind. For example, in meteorological applications data-assimilation methods like 4D-Var (e.g., Courtier et al., 1994) or the ensemble Kalman filter (Evensen, 1994) are employed. These methods, however successful, are limited in the sense that the analysis is linearised and, thus, the methods assume Gaussian distributions.

There exists an ensemble-based data-assimilation method that does not make such an assumption. It is particle filtering. In particle filtering, the probability distribution function of the state is approximated by an ensemble of particles, where a particle (or ensemble member) is a full model state obtained by running a model. In order to have non-identical particles a perturbation is applied to initial conditions, for example. Then, each particle is propagated forward in time using the model. When the observation becomes available, the so-called importance weights are assigned to the particles based on how close to the observation they are. Small weights are given to particles far from the observation; large weights, to particles close to the observation. The ensemble

mean, which is the best estimate of the state, is then a sum of the particles each multiplied by the corresponding weight.

Particle filtering has no assumption of gaussianity, uses a full nonlinear model to propagate the particles, but unfortunately, suffers from the "curse of dimensionality" (Snyder et al., 2008), meaning that for a high-dimensional system particles (ensemble members) tend to drift apart during their forward evolution leading, consequently, to large variance in the corresponding importance weights. If the ensemble size is small, after a few data-assimilation cycles all but one of the particles have importance weights close to zero, and an ensemble that has collapsed to a single particle can no longer approximate the probability distribution function of the state. Therefore, particle filtering has not yet been employed for operational geophysical problems. To overcome the limitation of degeneracy, a new particle filter has been introduced by van Leeuwen (2010), the equivalent weights particle filter. There, the particles are guided towards the observations during the model simulations and smaller variance in the particles weights is induced. The equivalent weights particle filter has shown good performance for the Lorenz-63 and the Lorenz-95 models (van Leeuwen, 2010), and for the barotropic vorticity equation (van Leeuwen and Ades, 2012).

Paleoclimate applications are somewhat different than meteorological applications. The system is nonlinear and high-dimensional as well, but the observations are sparse and have large uncertainties. Moreover, the available observations allow reconstructions of only large-scaled features averaged over several months or even years rather than a few tenths of kilometres and six hours scales. Therefore, for a paleoclimate application the number of degrees of freedom of the system can be reduced by performing spatial and temporal averages without substantial loss of needed information. This allows the use of a particle filter even without the "guidance" of van Leeuwen (2010). For instance, Goosse et al. (2009) used a particle filter with 96 members and the dataset HADCRUT3 (Brohan et al., 2006) to reconstruct the past half-century climate state in the Southern Hemisphere. It was shown that variables like surface air temperature averaged over large domains, sea ice area in the Southern Ocean, and the southern annual mode are in agreement with the observations at an annual time scale. Annan and Hargreaves (2012) assessed reconstructions of annual mean temperature anomalies over the Northern Hemisphere for the past two millennia. The reconstructions were obtained using a particle filter with 100 members and limited pseudo-proxies of surface temperature. It was pointed out that annual temperature at the hemispheric scale is well reconstructed, even when only 50 pseudo-proxies are used, as to the regional scale the performance is poor giving negative skill for the spatial field in some regions.

While those applications were dealing with annual reconstructions at a large spatial scale, our goal is to test data-assimilation methods to reconstruct the climate state with a more detailed spatial structure and on a seasonal time scale. Since the number of degrees of freedom is larger when estimating seasonal variability than when estimating annual variability and some improvements of particle filtering were introduced in van Leeuwen (2010) to handle high-dimensional systems with many degrees of freedom, we test this data-assimilation method, the nudging proposal particle filter, in our experiments. Moreover, we compare this method to a particle filter with sequential importance resampling – a method used in Dubinkina et al. (2011) – and to a nudging – a data-assimilation method widely used by general circulation models for initialising climate predictions (e.g., Keenlyside et al., 2008; Pohlmann et al., 2009; Swingedouw et al., 2012).

In our studies, we focus on the Southern Hemisphere as it is an interesting test case for the model dynamics that includes potentially complex interactions with sea ice. We employ the climate model of intermediate complexity LOVECLIM – a coupled model with atmospheric, oceanic and sea-ice components. As the period of interest we choose 150 yr from the year 1850 until 2000. Variation of the anthropogenic impact in 1850–2000 allows us to assess the performance of a data-assimilation method under different magnitudes of the forcing. Moreover, a study over a rather short time period for a paleoclimatological application gives the basis for future applications over longer periods.

Experiments with pseudo-proxies, which are derived from simulations of climate models, are quite typical for paleoclimatological applications (e.g., Smerdon, 2012), as they give more freedom in estimating skill of a method used to obtain a climate state reconstruction. Therefore, we constrain the model by pseudo-observations instead of instrumental records. We use pseudo-observations of surface air temperature anomalies, since for the last centuries observations of surface air temperature (either instrumental or proxy reconstructions) appear to be the most disposable. We perform two series of experiments: using the pseudo-observations given at every grid cell over the assimilated domain and using the pseudo-observations given at the same locations as the dataset of instrumental surface temperature records HADCRUT3 (Brohan et al., 2006). With the latter series we aim to approach a more realistic setup for a paleoclimatological application, but without leaving the twin-experiment framework. The design of our experiments is close to a real application of a data-assimilation method (e.g., Goosse et al., 2012) and, therefore, can be easily adapted for such an application. For a more distant past, the number of proxies is substantially smaller. Therefore, the performance of data assimilation is expected to be weaker, but still satisfactory, if the signal recored in sparse proxies is strong and the aim is to reconstruct large-scale features, e.g., Annan and Hargreaves (2012); Mathiot et al. (2013)

The paper is organised as follows: in Sect. 2, we give a description of three data-assimilation methods that are used for the past climate state reconstructions: the sequential importance resampling filter, a nudging and the nudging proposal particle filter. In Sect. 3, we describe the climate model

LOVECLIM and the experimental setup. Results of the experiments using the dense net of the pseudo-observations are given in Sect. 4. In Sect. 5, the performance of a data-assimilation method is addressed when the sparse net of the pseudo-observations is employed. Finally, conclusions are given in Sect. 6.

2 Data assimilation methods

2.1 Particle filter with sequential importance resampling

If the discrete equation for estimating the state ψ of a model at a time t_n, where n is the time index, is a function f of the state ψ at a time t_{n-1}

$$\psi^n = f(\psi^{n-1}), \tag{1}$$

then its M realisations, called particles, which are obtained using different initial conditions, determine an ensemble $\{\psi_i^n\}_{i=1}^M$ that represents the model probability density as following

$$p(\psi^n) = K^{-1}\sum_{i=1}^{M}\delta(\psi^n - \psi_i^n), \tag{2}$$

where δ is a kernel density and K (equal to M here) is a normalisation factor (Hereinafter any normalisation factor will be denoted by K. Since K is the same for every particle, it is irrelevant for the weight comparison). Given the observation d^n of the model state ψ^n Bayes theorem indicates that the posterior probability is

$$p(\psi^n|d^n) = K^{-1}p(d^n|\psi^n)p(\psi^n), \tag{3}$$

where $p(d^n|\psi^n)$ is the likelihood of the observations given the model state and it is related to the observation uncertainties. After substituting the density from Eq. (2) into Eq. (3), the posterior probability becomes

$$p(\psi^n|d^n) = \sum_{i=1}^{M} w_i^n \delta(\psi^n - \psi_i^n)$$
$$\text{with } w_i^n = K^{-1}p(d^n|\psi_i^n).$$

The weights $\{w_i^n\}_{i=1}^M$ are computed assuming that the likelihood $p(d^n|\psi_i^n)$ is Gaussian

$$p(d^n|\psi_i^n) = K^{-1}\exp$$
$$\left[-\frac{1}{2}(d^n - H(\psi_i^n))^T R^{-1}(d^n - H(\psi_i^n))\right]. \tag{4}$$

Here H is the measurement operator that projects a model state ψ_i^n to the location of the observation d^n, and R is the error covariance of the observations.

The last step of the sequential importance resampling filter consists of particles resampling according to their weights, which is often necessary to avoid the filter degeneracy. The particles with small weights are eliminated, whereas the particles with large weights are kept. To retain the total number of the particles the remaining particles are duplicated and perturbed. The perturbation is computed from the empirical orthogonal function analysis of the model error, which is the difference between a control model run and the instrumental records of surface temperature HADCRUT3 over the last 150 yr. The perturbation, which is a sum of the first ten modes each multiplied by a random scalar, is added to initial conditions for surface temperature. Then, the particles are propagated forward in time by the model until the next observation is available. After that, the importance of weights are computed again, but using the new observation, and the whole procedure is repeated until the end of the period of interest. For a more detailed description of the sequential importance resampling filter, we refer the reader to van Leeuwen (2009).

2.2 Nudging

Nudging consists of a term that is added to the prognostic model equation in order to pull the model state towards the observation (e.g., Hoke and Anthes, 1976). In a discrete form, we have

$$\psi^n = f(\psi^{n-1}) + \alpha H^T(d^n - H(\psi^{n-1})) + \xi^n, \tag{5}$$

where α is a nudging parameter, ξ^n is a stochastic noise and d^n is the observation. Presence of the additive stochastic noise ξ^n is not generally required for the nudging formulation but is, however, essential for the nudging proposal particle filter as it will follow later. In complex nudging schemes, the parameter α is a matrix that incorporates information of error correlation between variables. We, however, consider α to be a scalar matrix for simplicity and also because a scalar α is still used in many recent studies using climate models, (e.g., Swingedouw et al., 2012). Its choice, based on physical constraints, defines the strength of the nudging: a strong nudging can yield to a wrong dynamics due to a fast convergence of the solution to the observation, whereas a weak nudging provides with the solution that is unconstrained by the observation. In our experiments, we nudge sea surface temperature, since in general circulation models nudging is usually performed over the ocean (e.g., Swingedouw et al., 2012).

2.3 The nudging proposal particle filter

In the nudging proposal particle filter, like in the particle filter of Sect. 2.1, the model probability density is represented by an ensemble of particles according to Eq. (2), and the Bayes theorem, Eq. (3), is used to derive the posterior probability. The model equation, however, is distinct from Eq. (1). Let the model equation have the stochastic model error denoted by $\hat{\xi}^n$, which is related to unknown parameters of the model,

for example. Then,

$$\psi^n = f(\psi^{n-1}) + \hat{\xi}^n, \tag{6}$$

and the transitional density $p(\psi^n|\psi^{n-1})$ is the density of $\hat{\xi}^n$ with mean $f(\psi^{n-1})$. Moreover, if the model equation has the nudging term like in Eq. (5), one can define the proposal transition density $q(\psi^n|\psi^{n-1}, d^n)$ as the probability density of ξ^n with mean $f(\psi^{n-1}) + \alpha H^T(d^n - H(\psi^{n-1}))$. Taking into account both the transitional density and the proposal transition density when deriving the posterior probability gives the following weights

$$W_i^n = K^{-1} p(d^n|\psi_i^n) \prod_{r=m+1}^{n} \frac{p(\psi_i^r|\psi_i^{r-1})}{q(\psi_i^r|\psi_i^{r-1}, d^n)}. \tag{7}$$

Here, index m is related to a time t_m at which the observation d^m – the observation previous to d^n – is available. Therefore, when computing the product in Eq. (7), all the model states $\{\psi_i^m, \psi_i^{m+1}, \ldots, \psi_i^n\}$ between the two consecutive observations d^m and d^n are taken into account. If the observations are as frequent as the model states then $m+1=n$, otherwise $m+1<n$. Note, that the devision in Eq. (7) does not lead to singularity since the support of the proposal transition density $q(\psi_i^r|\psi_i^{r-1}, d^n)$ is equal or wider than the one of the transitional density $p(\psi_i^r|\psi_i^{r-1})$ due to the presence of the stochastic noise ξ^r in the model Eq. (5). For computing the weights we take $p(d^n|\psi_i^n)$ to be equal to Eq. (4), the transition density to be equal to

$$p(\psi_i^r|\psi_i^{r-1}) = K^{-1} \exp \left[-\frac{1}{2}(\psi_i^r - f(\psi_i^{r-1}))^T C^{-1} (\psi_i^r - f(\psi_i^{r-1})) \right],$$

and the proposal transition density to be equal to

$$q(\psi_i^r|\psi_i^{r-1}, d^n) = K^{-1} \exp \left[-\frac{1}{2}(\xi_i^r)^T \Sigma^{-1} \xi_i^r \right].$$

The model error covariances C and Σ for simplicity are taken to be equal.

Since the nudging does not guarantee small variance in the particles and, consequently, in the importance weights when many degrees of freedom are present, the nudging proposal particle filter can still become degenerative. Therefore, the model states has to be adjusted just before the calculation of the weights such that the weights do not differ substantially afterwards, and the nudging proposal particle filter becomes the equivalent weight particle filter (van Leeuwen, 2010). We, however, leave out this part of "almost equal weights" since the number of degrees of freedom in our application is still quite small, and use the nudging proposal particle filter. For a comprehensive explanation of the nudging proposal particle filter the reader is referred to van Leeuwen (2010); van Leeuwen and Ades (2012).

3 Description of experimental setup

The three-dimensional Earth system model of intermediate complexity LOVECLIM1.2 (Goosse et al., 2010) used here consists of the atmospheric component ECBILT2 (Opsteegh et al., 1998), the oceanic component CLIO3 (Goosse and Fichefet, 1999), and the terrestrial vegetation module VECODE (Brovkin et al., 2002). The atmospheric model is a three level quasi-geostrophic model of horizontal resolution T21 that includes a radiative scheme and a parametrisation of the heat exchanges with the surface. The free-surface ocean model is an ocean general circulation model coupled to a sea-ice model with horizontal resolution of three by three degrees and 20 unevenly spaced vertical levels in the ocean. The vegetation module describes annual changes in vegetation cover taking into account trees, grass and deserts; its horizontal resolution matches the resolution of the atmospheric component.

Starting from an equilibrium run with LOVECLIM1.2 in the year 850, we perform a transient simulation until 2000, from which we derive the pseudo-observations for 1850–2000. Four additional simulations over the period 850–1850 using perturbed initial conditions provide with the conditions in the year 1850 used to initialise the data-assimilation experiments. These experiments are constrained by the pseudo-observations of surface air temperature described above, to which we add a Gaussian noise with standard deviation 0.5 °C in order to mimic the instrumental error. When comparing the reconstructions with the pseudo-observations no noise, however, is applied to the pseudo-observations, meaning that the comparison is done with the truth.

In the data-assimilation experiments the particle filters (either with sequential importance resampling or the nudging proposal one), we use the pseudo-observations averaged on a seasonal scale. The seasonal scale is small enough to provide with detailed climate state reconstructions, but large enough not to impose the issue of degeneracy of the particle filters. Moreover, we apply a spatial filter to the particles as in Dubinkina et al. (2011) before computing the importance of weights in order to reduce the number of degrees of freedom.

In the nudging (either alone or as a part of the nudging proposal particle filter), we use the pseudo-observations of monthly mean surface air temperature, since the nudging does not degenerate and monthly averages are generally the smallest scales of observations in long-term applications. The nudging is performed over the global ocean by introducing a term into the computation of heat fluxes between the atmosphere and the ocean. The nudging parameter α is chosen such that a corresponding relaxing timescale for a mixed layer of 50 m depth is 20 days. Moreover, a maximum of 50 W m^{-2} is imposed on the heat flux adjustment due to the nudging. This nudging is 3 times stronger than the nudging used in Swingedouw et al. (2012) and 5 times weaker than the nudging used in Keenlyside et al. (2008). The stochastic error ξ^n is constructed as following: we perform empirical

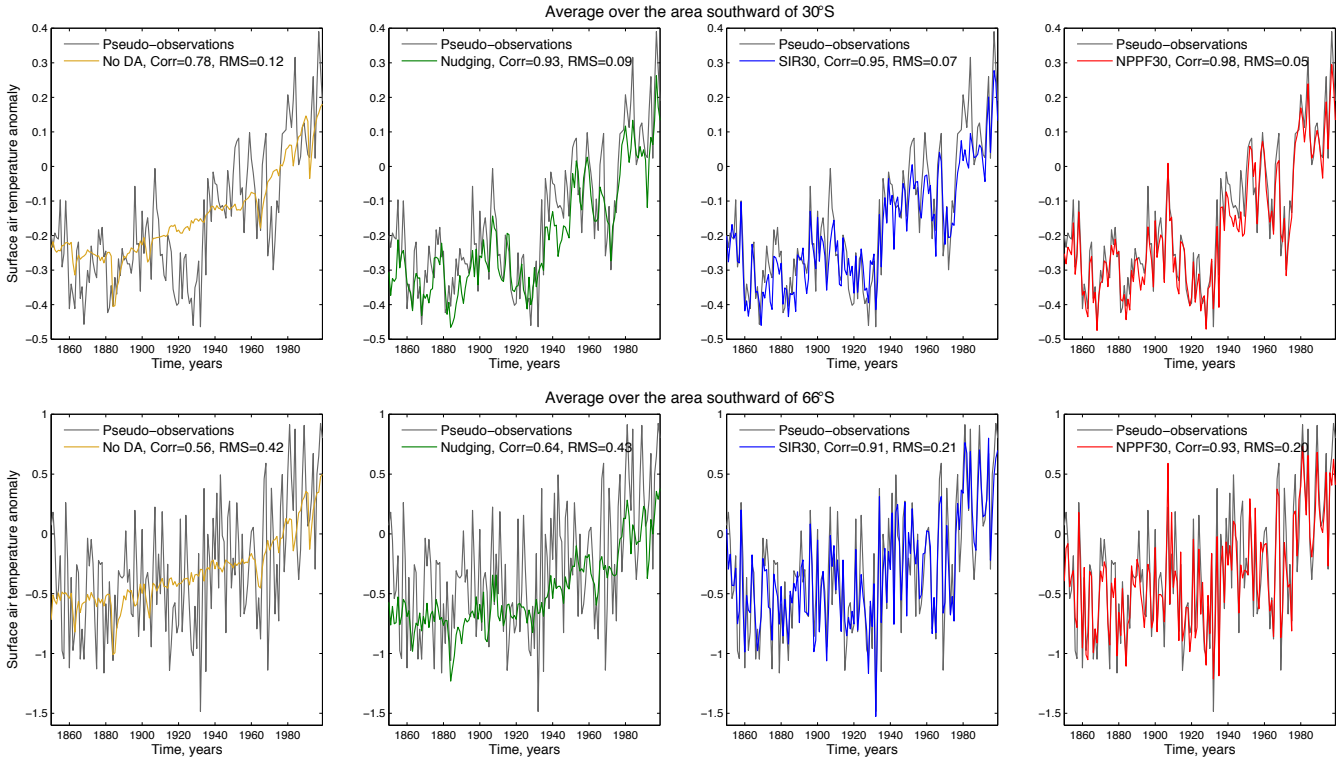

Fig. 1. Reconstructions of surface air temperature anomalies (°C) averaged over the area southward of 30° S (top panel) and over the area southward of 66° S (bottom panel), when the dense pseudo-observations are assimilated. Gray line: pseudo-observations; yellow line: model simulations without data assimilation; green line: the nudging; blue line: the sequential importance resampling applied over the polar cap southward of 30° S; red line: the nudging proposal particle filter applied over the polar cap southward of 30° S. Correlations and the RMS errors are displayed in upper left corners.

orthogonal function (EOF) analysis of the model error, which is the difference between the output of a control model run and instrumental surface temperature records HADCRUT3 (Brohan et al., 2006) over the last 150 yr. Then, the noise is a sum of the first ten modes each multiplied by a random coefficient and this noise, together with the nudging term, is added to the equation of heat fluxes.

In all three data-assimilation methods, we use 96 particles, which seems to be sufficient for representing the probability density and is computationally affordable. The error covariance of the observations R is computed using the instrumental error and the error of representativeness, as in Dubinkina et al. (2011), and the model error covariance C is assumed to be a scalar matrix with $(0.5\,°\mathrm{C})^2$ on the diagonal. The latter assumption is relatively crude and it would be more adequate to build the covariance matrix C by taking into account spatial correlation of the model error. We, however, consider a scalar matrix for representing the covariance matrix C and a scalar matrix for the nudging in order to assess the performance of the nudging proposal particle filter when it was obtained by a transition density of a simple form.

4 Assimilation of dense pseudo-observations

In the following experiments, the pseudo-observations of surface air temperature are given at every grid cell. Since assimilation of the pseudo-observations over the whole globe leads to filter degeneracy and assimilation over a small domain does not take many pseudo-observations into account, we make a compromise by choosing an area covering the polar cap southward of 30° S. Nudging is, however, applied over the global ocean, whether it is a part of the nudging proposal particle filter or the nudging method itself.

We examine the reconstructions of surface air temperature averaged over two domains: the area southward of 30° S and the area southward of 66° S (the top and bottom panels of Fig. 1). To assess the performance of a method, we compute correlation and the root-mean-square (RMS) error between reconstructed surface air temperature and the pseudo-observations. For the area southward of 30° S, shown in the top panel of Fig. 1, all three data-assimilation methods perform well giving the low RMS errors and high correlations. Data assimilation contributes considerably to these high correlations, since for simulations without data assimilation correlation due to the response to common forcing is only 0.78.

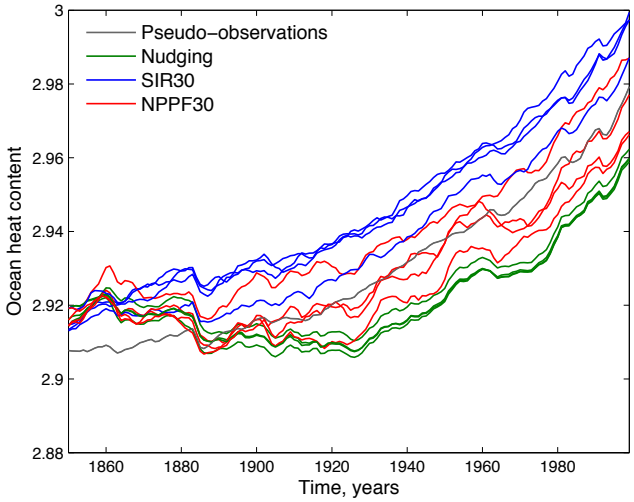

Fig. 2. Reconstructions of ocean heat content (°C), when the dense pseudo-observations are assimilated. Gray line: pseudo-observations; green lines: the nudging; blue lines: the sequential importance resampling filter applied over the polar cap southward of 30° S; red lines: the nudging proposal particle filter applied over the polar cap southward of 30° S.

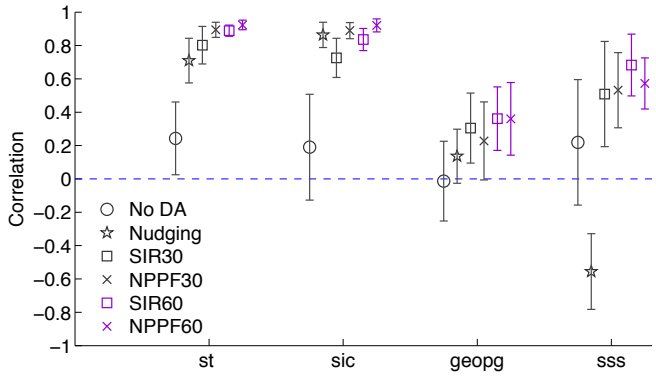

Fig. 3. Assimilation of dense pseudo-observations. Correlations between first PCs of the pseudo-observations and projections of the model simulations onto the corresponding first EOFs of the pseudo-observations for different variables: st is for surface temperature, sic is for sea ice concentration, geopg is for geopotential height, sss is for sea surface salinity. EOFs are computed for May–October of twenty-four 21-yr periods over the area southward of 60° S. The circle is the mean correlation for simulations without data assimilation; the star is the mean correlation for the model simulations using the nudging; the gray square is the mean correlation for the model simulations using the sequential importance resampling filter over the polar cap southward of 30° S; the gray cross is the mean correlation for the model simulations using the nudging proposal particle filter over the polar cap southward of 30° S; the purple square is the mean correlation for the model simulations using the sequential importance resampling filter over the polar cap southward of 60° S; the purple cross is the mean correlation for the model simulations using the nudging proposal particle filter over the polar cap southward of 60° S. Error bars correspond to one standard deviation.

For the area southward of 66° S, shown in the bottom panel of Fig. 1, the performance of the particle filters is only a bit weaker than for the area southward of 30° S unlike in the case of the nudging, where correlation decreases from 0.93 to 0.64, the RMS error increases from 0.09 to 0.43 °C, and the variance of the reconstructed anomaly (green curve) is smaller than the variance of the pseudo-observations (gray curve). Weaker performance of the nudging for the area southward of 66° S is due to the fact that the ocean covers a small fraction of the surface southward of 66° S; therefore, since the nudging is done over the ocean only, it has a weaker direct influence on this area and propagation of the signal from the ocean to the land is not strong enough to lead to high correlations. Nevertheless, the nudging still outperforms the simulations without data assimilation, for which the correlation is smaller (0.56).

For estimating the performance of a data-assimilation method for the ocean reconstruction we consider ocean heat content, which is represented by mean ocean temperature in LOVECLIM. We perform four experiments using different initial conditions for each data-assimilation method. Figure 2 illustrates that ocean heat content is not significantly altered by the sequential importance resampling filter (blue curves). Therefore, the sequential importance resampling filter does not change the heat budget of the climate model. Since the same forcing is used for deriving the pseudo-observations and when performing the data-assimilation experiments, ocean heat content from the sequential importance resampling filter is parallel to the pseudo-observations reflecting the influence of different initial conditions during the whole period. The nudging, on the contrary, has a strong influence on ocean heat content (green curves). This is due to the way the nudging is implemented: it adjusts heat fluxes from the atmosphere to the ocean. Consequently, ocean temperature changes, so does ocean heat content. The nudging proposal particle filter obtains ocean heat contents (red curves) that appear to be the closest to the pseudo-observations (gray curve) and has the smallest mean RMS error over four experiments (0.008 °C against 0.014 °C for the sequential importance resampling filter and 0.013 °C for the nudging).

Next, we investigate the skill of the assimilation methods in reconstructing spatial features. In order to do that, we compute first empirical orthogonal functions (EOFs) of the pseudo-observations and project the results of model simulations onto them. Then, the corresponding principal components (PCs) and the projections are compared by means of correlation. We perform four experiments using different initial conditions for every data-assimilation method. The EOFs are computed for winter period (from May until October) over the area southward of 60° S, as we are mainly interested in the regions that are ice covered or that are close to the ice edge, and over a 21-yr period, since it is long enough to capture the main features of the state by the EOFs and short

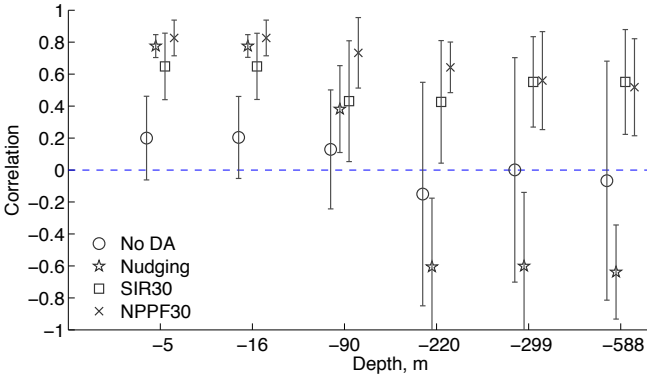

Fig. 4. Assimilation of dense pseudo-observations. Correlations between first PCs of the pseudo-observations and projections of the model simulations onto the corresponding first EOFs of the pseudo-observations for ocean temperature at different depths. EOFs are computed for May–October of six 21-yr periods over the area southward of 60° S. The circle is the mean correlation for simulations without data assimilation; the star is the mean correlation for the model simulations using the nudging; the square is the mean correlation for the model simulations using the sequential importance resampling filter over the polar cap southward of 30° S; the cross is the mean correlation for the model simulations using the nudging proposal particle filter over the polar cap southward of 30° S. Error bars correspond to one standard deviation.

enough to split one model run in several such periods. Therefore, we divide a 150-yr run in six 21-yr periods starting from the year 1865 and ending in the year 1990, skipping the first 15 yr to avoid the bias induced by the initial conditions. Performing the EOF analysis over six 21-yr periods from four different experiments gives twenty-four correlations for every data-assimilation method.

In Fig. 3, we plot mean correlations plus and minus one standard deviation for different variables and different data-assimilation methods. When reconstructing surface air temperature (st) or sea ice concentration (sic) all three methods perform rather well resulting in high correlations, as shown in gray colours in Fig. 3. For these variables, the simulations without data assimilation (circles) give low but positive correlations, which is due to the employment of the same forcing when deriving the pseudo-observations. The skill of the nudging proposal particle filter (gray cross) when reconstructing surface air temperature (st) is only slightly higher than the skill of the sequential importance resampling filter (gray square). When reconstructing sea ice concentration (sic), the nudging proposal particle filter (gray cross) shows evident improvement compared to the sequential importance resampling filter (gray square). Here, the nudging and the nudging proposal particle filter have the higher skills than the sequential importance resampling filter, since the corrections of heat fluxes from the atmosphere to the ocean have a strong impact on sea ice concentration.

Changes in atmospheric circulation is an important characteristics of past climate variability in the Southern Ocean (e.g., Lefebvre and Goosse, 2008; Yuan and Li, 2008). Pressure observations that can be used to constrain the model in order to get reliable estimations of atmospheric circulation are, however, very limited for paleoclimate applications. Therefore, we investigate the skill of the atmospheric circulation reconstructions when surface air temperature is assimilated. We perform the EOF analysis for geopotential height, the variable in LOVECLIM that represents atmospheric circulation. The correlations for geopotential height (geopg) are shown in gray colours in Fig. 3. The correlations of simulations with data assimilation are lower than for surface temperature (st) or sea ice concentration (sic), as expected, but still overall positive. Moreover, their mean correlations (gray star, square and cross) are higher than the mean correlation of simulations without data assimilation (circle), meaning that data assimilation of surface air temperature improves the reconstruction of atmospheric circulation – a variable that is not directly linked to the pseudo-observations.

Next, we perform the EOF analysis of sea surface salinity (sss), whose variations play a crucial role in the changes in ocean density and, consequently, in the oceanic circulation and the vertical stability of the water column (e.g., Martinson, 1990; Gordon, 1991). In Fig. 3, we see that the nudging proposal particle filter (gray cross) and the sequential importance resampling filter (gray square) provide with positive and rather good correlations, taken into account that sea surface salinity is not directly linked to assimilated surface air temperature. By contrast, sea surface salinity obtained by the nudging (gray star) has always negative correlations with the pseudo-observations. This could be due to the strong nudging used in this study, since Swingedouw et al. (2012) argued that a short relaxing timescale of 4 days used in Keenlyside et al. (2008) resulted in non-physical heat fluxes, while a relaxing timescale of 60 days used in Swingedouw et al. (2012) did not impose such a spurious behaviour of ocean heat content. Therefore, we perform experiments with a weaker nudging with a relaxing timescale of 60 days, as in Swingedouw et al. (2012). This nudging is 3 times weaker than the nudging used in our standard experiment. Correlations between sea surface salinity obtained by the weaker nudging and the pseudo-observations are still negative (not shown). Hence, our results do not appear to depend crucially on the value of a relaxing timescale. In order to understand the reason why the pattern of sea surface salinity obtained by the nudging is opposite to the pseudo-observations, we perform the EOF analysis for ocean temperature at different depths. As it is seen in Fig. 4, the nudging adjusts ocean temperature near the surface, but does not respect the dynamics of the ocean. In particular, the nudging term strongly modifies the mixing (not shown) leading to a wrong vertical ocean temperature profile and to wrong vertical salinity. Thus in climate models, nudging has to be used with caution at least when applied to a region like the Southern Ocean where a small modification of

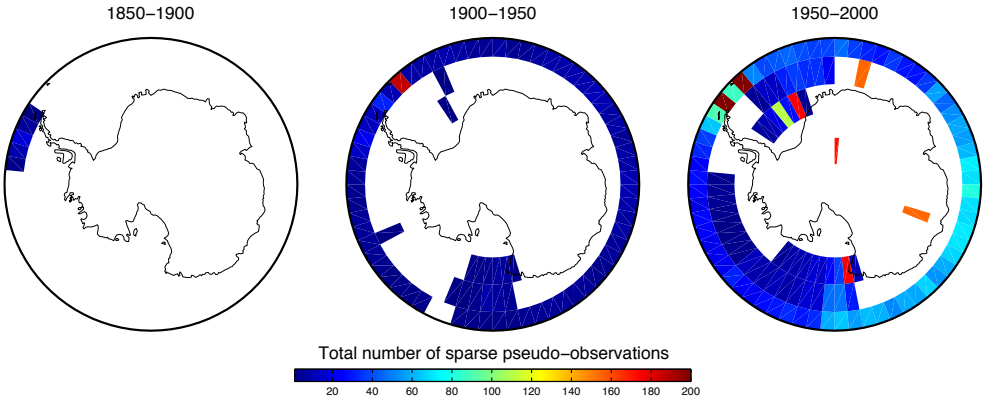

Fig. 5. Number of seasons of the sparse pseudo-observations per grid cell southward of 60° S over different time periods: 1850–1900 (left), 1900–1950 (middle), and 1950–2000 (right).

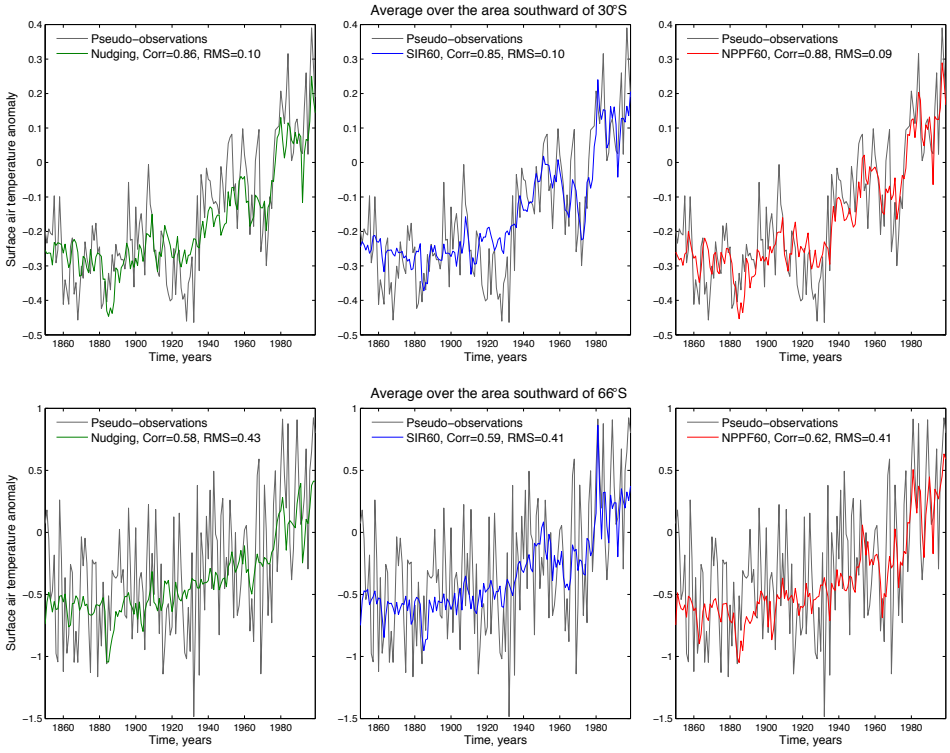

Fig. 6. Reconstructions of surface air temperature anomalies (°C) averaged over the area southward of 30° S (top panel) and over the area southward of 66° S (bottom panel), when the sparse pseudo-observations are assimilated. Gray line: pseudo-observations; green line: the nudging; blue line: the sequential importance resampling applied over the polar cap southward of 60° S; red line: the nudging proposal particle filter applied over the polar cap southward of 60° S. Correlations and the RMS errors are displayed in upper left corners.

the vertical stratification of the ocean variables has a strong impact on the surface.

Next, we investigate whether it is possible to increase correlations by assimilating the pseudo-observations over a smaller domain – the area southward of 60° S. As shown in Fig. 3, correlations for all variables are higher when the assimilation domain is smaller – the purple squares compared to the gray ones for the sequential importance resampling filter and the purple crosses compared to the gray ones for the nudging proposal particle filter. Keeping the total number of particles the same (96), but reducing the assimilation area, thus, reducing the number of degrees of freedom, results in more particles with relatively high importance weights. Consequently, the mean obtained by these particles has higher correlation and the smaller RMS error with the pseudo-observations.

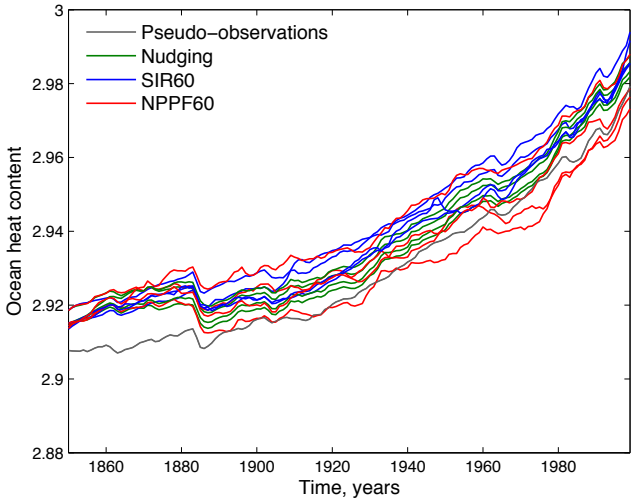

Fig. 7. Reconstructions of ocean heat content (°C), when the sparse pseudo-observations are assimilated. Gray line: pseudo-observations; green lines: the nudging; blue lines: the sequential importance resampling filter applied over the polar cap southward of 60° S; red lines: the nudging proposal particle filter applied over the polar cap southward of 60° S.

5 Assimilation of sparse pseudo-observations

In the following experiments, we investigate the performance of the data-assimilation methods when the pseudo-observations are as sparse as the dataset HADCRUT3 of the instrumental surface temperature records over the last 150 yr (Brohan et al., 2006) by selecting the pseudo-observations at the same locations as the HADCRUT3 dataset. The spatial resolution of the sparse pseudo-observations changes in time as it can be seen in Fig. 5, where number of seasons of the sparse pseudo-observations per grid cell southward of 60° S are shown over different time periods: 1850–1900 (the left panel), 1900–1950 (the middle panel), and 1950–2000 (the right panel). We assimilate the sparse pseudo-observations over the area southward of 60° S in order to decrease the number of degrees of freedom and avoid degeneracy. The nudging is still applied over the global ocean, but at the HADCRUT3 locations only.

In Fig. 6, we plot time series of surface air temperature anomalies averaged over the area southward of 30° S (the top panel) and over the area southward of 66° S (the bottom panel). Compared to the case of assimilating the dense pseudo-observations, the variance of the anomalies is underestimated, which is due to the sparsity of the pseudo-observations. Annan and Hargreaves (2012) have also observed the decrease in the variance when the pseudo-observations become more sparse. Nevertheless, all three data-assimilation methods estimate still reasonably well surface air temperature averaged over the area southward of 30° S: correlations are 0.86, 0.85 and 0.88 for the nudging, the sequential importance resampling filter, and the nudging proposal particle filter, respectively, while for the model without any data assimilation correlation is 0.78. In the area southward of 66° S where only a few pseudo-observations are located, we have a good estimation of the trend, but not of the variance. Moreover, the trend reconstruction is achieved mainly due to the well-defined forcing, not due to data assimilation. Indeed, when no data assimilation is used correlation is 0.56 and correlations obtained by the data-assimilation methods are 0.58, 0.59 and 0.62 for the nudging, the sequential importance resampling filter, and the nudging proposal particle filter, respectively. It should be mentioned, however, that when the forcing is unknown and a random one is applied, the trend as well as the forcing can be still estimated due to data assimilation, see Dubinkina et al. (2011).

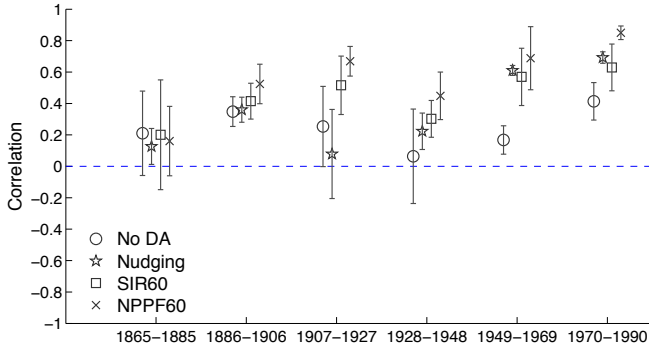

Fig. 8. Assimilation of sparse pseudo-observations. Correlations between first PCs of the pseudo-observations and projections of the model simulations onto the corresponding first EOFs of the pseudo-observations for surface air temperature for different time periods. EOFs are computed over the area southward of 60° S for May–October of 21-yr periods from four runs. The circle is the mean correlation for simulations without data assimilation; the star is the mean correlation for the model simulations using the nudging; the square is the mean correlation for the model simulations using the sequential importance resampling filter over the polar cap southward of 60° S; the cross is the mean correlation for the model simulations using the nudging proposal particle filter over the polar cap southward of 60° S. Error bars correspond to one standard deviation.

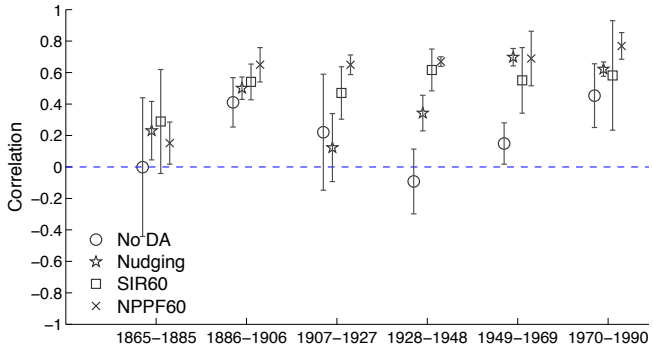

Fig. 9. Same as Fig. 8, but for sea ice concentration.

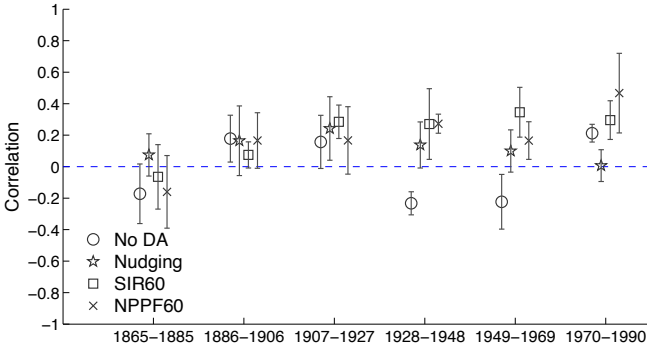

Fig. 10. Same as Fig. 8, but for geopotential height.

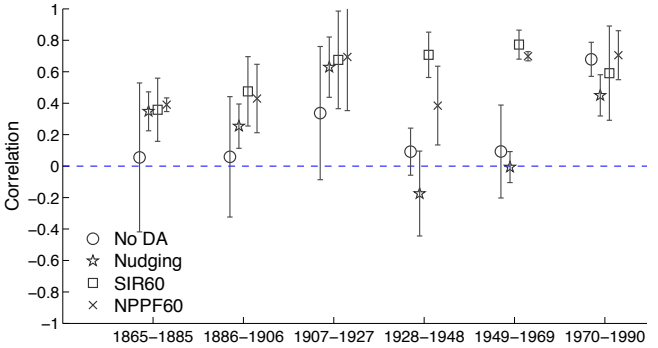

Fig. 11. Same as Fig. 8, but for sea surface salinity.

In Fig. 7, we plot ocean heat content obtained by the data-assimilation methods. We use four different initial conditions, which resulted in four reconstructions per method. Overall all three methods overestimate the truth (gray line), unlike in the case of assimilating the dense pseudo-observations, when the sequential importance resampling filter overestimates it, the nudging underestimates and the nudging proposal particle filter provides with the smallest mean RMS error. Here, the mean RMS errors are comparable: 0.008 °C for the nudging, 0.009 °C for the nudging proposal particle filter, and 0.01 °C for the sequential importance resampling filter.

For examining the spatial skill of the data-assimilation methods, we perform the EOF analysis as described in Sect. 4, but since resolution of the sparse pseudo-observations depends on time, the six 21-yr periods are not equivalent anymore. Therefore, we perform four runs for every data-assimilation method using different initial conditions and compare every 21-yr period separately. The comparison is done by computing correlations between first PCs of the dense pseudo-observations and projections of model simulations onto the corresponding first EOFs of the pseudo-observations. Mean and standard deviation are computed over four correlations.

Figure 8 illustrates that the nudging proposal particle filter provides overall with higher correlations than any other method when reconstructing surface air temperature. Compared to the case of assimilating the dense pseudo-observations, assimilation of the sparse pseudo-observations results in smaller mean correlations and larger standard deviations, except for the period 1970–1990, at which many pseudo-observations are available.

In Fig. 9, we see that starting from the period 1886–1906 correlations given by the nudging proposal particle filter and by the sequential importance resampling filter are quite reasonable for sea ice concentration. Moreover, the mean correlations of the nudging proposal particle filter are higher than the ones of the sequential importance resampling filter, just as for surface air temperature correlations shown in Fig. 8. The correlations given by the nudging for the periods 1907–1927 and 1928–1948 are significantly worse than the ones given by the particle filters. For 1907–1927 the mean correlation of the nudging is even smaller than the mean correlation of simulations without data assimilation, unlike in the case of assimilating the dense pseudo-observations, when for every time period the nudging provides with the mean correlation higher than the mean correlation of simulations without data assimilation (not shown). Therefore, when the pseudo-observations have low density the nudging may not be able to propagate the sparse signal.

From Fig. 10, we see that the data-assimilation methods do not constrain the model well enough in order to have high correlations for atmospheric circulation. Only in the period 1970–1990 with many pseudo-observations, correlations improve and become comparable to the correlations for geopotential height when assimilating the dense pseudo-observations over the area southward of 60° S (shown in purple in Fig. 3).

When reconstructing sea surface salinity, which is displayed in Fig. 11, the particle filters perform quite well: their mean correlations are substantially higher than the mean correlations of simulations without data assimilation, except for the period 1970–1990, when the mean correlation of the sequential importance resampling filter is a bit lower. As in the case of assimilating the dense pseudo-observations, the nudging obtains the sea surface salinity patterns opposite to the pseudo-observations, but only over some periods in this case (1928–1948 and 1949–1969).

6 Conclusions

We have shown that the nudging proposal particle filter provides with encouraging results: global variables like ocean heat content and surface air temperature averaged over large domains are well estimated. When assimilating the dense pseudo observations, the nudging proposal particle filter provides with reasonable reconstructions of not only variables that are directly linked to the pseudo-observations such as surface air temperature and sea ice concentration, but also variables such as geopotential height and sea surface salinity.

Reliable reconstructions of the latter variables are essential for paleoclimate applications since the observations of pressure and salinity are limited there. Moreover, these reconstructions give good perspectives for initialising climate predictions.

When assimilating the sparse pseudo-observations that are given at the same locations as the dataset of instrumental surface temperature records HADCRUT3, the performance of the nudging proposal particle filter is weaker due to the limited number of the pseudo-observations. Nevertheless, even at the end of the 19th century, the reconstructions of surface air temperature and of sea ice concentration have relatively high correlations. The reconstructions of geopotential height and of sea surface salinity have, however, lower correlations.

Overall, the nudging proposal particle filter achieves better or equivalent results compared to the sequential importance resampling filter. To be more precise, surface air temperature reconstructed by the sequential importance resampling filter has already high correlations with the pseudo-observations, and the nudging proposal particle filter introduces only a slight improvement. In reconstructing sea ice concentration, however, a clear improvement is accomplished by the nudging proposal particle filter compared to the sequential importance resampling filter. This efficiency of the nudging proposal particle filter is due to the small ensemble size, since with an infinitely large ensemble both of the particle filters will converge to the same posterior probability density function (van Leeuwen, 2010). When reconstructing variables that are not directly linked to the pseudo-observations such as geopotential height and sea surface salinity, the performance of the nudging proposal particle filter is equivalent to the performance of the sequential importance resampling filter.

When the pseudo-observations are dense, the nudging used here provides with good reconstructions of surface air temperature and of sea ice concentration, but it has the drawback of not respecting the dynamics of the ocean, which results in the wrong vertical profile of ocean temperature and, consequently, in wrong pattern of sea surface salinity. Nudging of salinity and ocean temperature over the depth could be a possible solution to this problem, but it should be kept in mind that observations of deep ocean temperature and salinity are available only for the recent past. Additionally, when the pseudo-observations are sparse, the nudging fails to propagate the signal from those sparse pseudo-observations.

The improvement brought by the nudging proposal particle filter is apparent, which makes a strong argument for the use of the nudging proposal particle filter in the climate state estimations. Some developments, however, are still needed in order to get better estimations of variables that are not strongly linked through the model dynamics to the assimilated surface air temperature such as geopotential height and salinity. Therefore, a more complex nudging (non-scalar matrix α) and better approximations of covariance matrices C and R should be considered, as it was discussed in van Leeuwen (2010).

Acknowledgements. H. Goosse is Senior Research Associate with the F. R. S. – FNRS-Belgium. The research leading to these results has received funding from the F. R. S. – FNRS, the Belgian Federal Science Policy Office and the European Union's Seventh Framework programme (FP7/2007-2013) under grant agreement no. 243908, "Past4Future. Climate change – Learning from the past climate". This is Past4Future contribution no. 44.

Edited by: S. Bronnimann

References

Annan, J. and Hargreaves, J.: Identification of climatic state with limited proxy data, Clim. Past, 8, 1141–1151, doi:10.5194/cp-8-1141-2012, 2012.

Bhend, J., Franke, J., Folini, D., Wild, M., and Brönnimann, S.: An ensemble-based approach to climate reconstructions, Clim. Past, 8, 963–976, doi:10.5194/cp-8-963-2012, 2012.

Brohan, P., Kennedy, J., Harris, I., Tett, S., and Jones, P.: Uncertainty estimates in regional and global observed temperature changes: a new dataset from 1850, J. Geophys. Res., 111, D12106, doi:10.1029/2005JD006548, 2006.

Brovkin, V., Bendtsen, J., Claussen, M., Ganopolski, A., Kubatzki, C., Petoukhov, V., and Andreev, A.: Carbon cycle, vegetation and climate dynamics in the Holocene: experiments with the CLIMBER-2 model, Global. Biogeochem. Cy., 16, 1139, doi:10.1029/2001GB001662, 2002.

Courtier, P., Thépaut, J.-N., and Hollingsworth, A.: A strategy for operational implementation of 4D-VAR, using an incremental approach, Q. J. Roy. Meteor. Soc., 120, 1367–1387, doi:10.1002/qj.49712051912, 1994.

Dubinkina, S., Goosse, H., Sallaz-Damaz, Y., Crespin, E., and Crucifix, M.: Testing a particle filter to reconstruct climate changes over the past centuries, Int. J. Bifurcat. Chaos, 21, 3611–3618, doi:10.1142/S0218127411030763, 2011.

Evensen, G.: Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics, J. Geophys. Res., 99, 10143–10162, doi:10.1029/94JC00572, 1994.

Goosse, H. and Fichefet, T.: Importance of ice-ocean interactions for the global ocean circulation: a model study, J. Geophys. Res., 104, 23337–23355, 1999.

Goosse, H., Lefebvre, W., de Montety, A., Crespin, E., and Orsi, A.: Consistent past half-century trends in the atmosphere, the sea ice and the ocean at high southern latitudes, Clim. Dynam., 33, 999–1016, doi:10.1007/s00382-008-0500-9, 2009.

Goosse, H., Brovkin, V., Fichefet, T., Haarsma, R., Huybrechts, P., Jongma, J., Mouchet, A., Selten, F., Barriat, P.-Y., Campin, J.-M., Deleersnijder, E., Driesschaert, E., Goelzer, H., Janssens, I., Loutre, M.-F., Morales Maqueda, M. A., Opsteegh, T., Mathieu, P.-P., Munhoven, G., Pettersson, E. J., Renssen, H., Roche, D. M., Schaeffer, M., Tartinville, B., Timmermann, A., and Weber, S. L.: Description of the Earth system model of intermediate complexity LOVECLIM version 1.2, Geosci. Model Dev., 3, 603–633, doi:10.5194/gmd-3-603-2010, 2010.

Goosse, H., Crespin, E., Dubinkina, S., Loutre, M., Mann, M., Renssen, H., Sallaz-Damaz, Y., and Shindell, D.: The role of forcing and internal dynamics in explaining the "Medieval Climate Anomaly", Clim. Dynam., 39, 2847–2866, doi:10.1007/s00382-012-1297-0, 2012.

Gordon, A.: Two Stable Modes of Southern Ocean Winter Stratification, in: Deep Convection and Deep Water Formation in the Oceans Proceedings of the International Monterey Colloquium on Deep Convection and Deep Water Formation in the Oceans, edited by: Chu, P. and Gascard, J., 57 of Elsevier Oceanography Series, 17–35, Elsevier, doi:10.1016/S0422-9894(08)70058-8, 1991.

Hoke, J. and Anthes, R.: The initialization of numerical models by a dynamic relaxation technique, Mon. Weather Rev., 104, 1551–1556, 1976.

Keenlyside, N., Latif, M., Jungclaus, J., Kornblueh, L., and Roeckner, E.: Advancing decadal-scale climate prediction in the North Atlantic sector, Nature, 53, 84–88, doi:10.1038/nature06921, 2008.

Lefebvre, W. and Goosse, H.: An analysis of the atmospheric processes driving the large-scale winter sea-ice variability in the Southern Ocean, J. Geophys. Res., 113, C02004, doi:10.1029/2006JC004032, 2008.

Martinson, D.: Evolution of the Southern Ocean winter mixed layer and sea ice: open ocean deepwater formation and ventilation, J. Geophys. Res., 95, 11641–11654, doi:10.1029/JC095iC07p11641, 1990.

Mathiot, P., Goosse, H., Crosta, X., Stenni, B., Braida, M., Renssen, H., Van Meerbeeck, C. J., Masson-Delmotte, V., Mairesse, A., and Dubinkina, S.: Using data assimilation to investigate the causes of Southern Hemisphere high latitude cooling from 10 to 8 ka BP, Clim. Past, 9, 887–901, doi:10.5194/cp-9-887-2013, 2013.

Opsteegh, J., Haarsma, R., Selten, F., and Kattenberg, A.: ECBILT: A dynamic alternative to mixed boundary conditions in ocean models, Tellus, 50A, 348–367, doi:10.1034/j.1600-0870.1998.t01-1-00007.x, 1998.

Pohlmann, H., Jungclaus, J., Köhl, A., Stammer, D., and Marotzke, J.: Initializing decadal climate predictions with the gecco oceanic synthesis: effects on the north atlantic, J. Climate, 22, 3926–3938, doi:10.1175/2009JCLI2535.1, 2009.

Smerdon, J.: Climate models as a test bed for climate reconstruction methods: pseudoproxy experiments, WIREs Clim. Change, 3, 63–77, doi:10.1002/wcc.149, 2012.

Snyder, C., Bengtsson, T., Bickel, P., and Anderson, J.: Obstacles to high-dimensional particle filtering, Mon. Weather Rev., 136, 4629–4640, doi:10.1175/2008MWR2529.1, 2008.

Swingedouw, D., Mignot, J., Labtoulle, S., Guilyardi, E., and Madec, G.: Initialisation and predictability of the AMOC over the last 50 years in a climate model, Clim. Dynam., 40, 2381–2399, doi:10.1007/s00382-012-1516-8, 2012.

Talagrand, O.: Assimilation of observations, an introduction, J. Metorol. Soc. JPN, Special Issue 75, 191–209, 1997.

van Leeuwen, P.: Particle Filtering in Geophysical Systems, Mon. Weather Rev., 137, 4089–4114, doi:10.1175/2009MWR2835.1, 2009.

van Leeuwen, P.: Nonlinear data assimilation in geosciences: an extremely efficientp article filter, Q. J. Roy. Meteor. Soc., 136, 1991–1999, doi:10.1002/qj.699, 2010.

van Leeuwen, P. and Ades, M.: Efficient fully nonlinear data assimilation for geophysical fluid dynamics, Comput. Geosci., 55, 16–27, doi:10.1016/j.cageo.2012.04.015, 2012.

Widmann, M., Goosse, H., van der Schrier, G., Schnur, R., and Barkmeijer, J.: Using data assimilation to study extratropical Northern Hemisphere climate over the last millennium, Clim. Past, 6, 627–644, doi:10.5194/cp-6-627-2010, 2010.

Yuan, X. and Li, C.: Climate modes in southern high latitudes and theur impact on Antarctic sea ice, J. Geophys. Res., 113, C06S91, doi:10.1029/2006JC004067, 2008.