The ECCO Report Series [1]

Efficiency of Reduced-Order Time-dependent Adjoint Data Assimilation Approaches

Ibrahim Hoteit, Armin Köhl, and and Detlef Stammer

Scripps Institution of Oceanography

Report Number **27**

Submitted to the

Ocean Modelling

January, 2004.

[1] The ECCO Project is a consortium of the Jet Propulsion Laboratory, The Massachusetts Institute of Technology and the Scripps Institution of Oceanography and is funded through a grant from the National Oceanographic Partnership Program (NOPP). Copies of this Report are available at www.ecco-group.org or from Detlef Stammer, Scripps Institution of Oceanography, La Jolla CA 92093-0230, ph.: (858) 822-3376; fax: (858) 534-4464; e-mail: dstammer@ucsd.edu.

Abstract

Adjoint data assimilation remains computationally demanding and limit thereby its application for operational use. To some extent this is due to the large dimensionality of the system control vector, i.e., the variables that are being adjusted to bring the model into consistency with the observations. To improve the convergence rate of an optimization, reduced-order optimization methods were suggested in the past. In this paper, we show that such order reduction applied to the control vector can indeed speed up the convergence of the optimization approach in a realistic setting. However, the order reduction can only speed up the initial convergence, since important features on smaller spatial scales can not be represented in the reduced space. A best performance was therefore obtained with a hybrid approach where we start with a reduced order control space but use the full control space subsequently. Numerical experiments show that the effectiveness of this strategy, while assimilating real data into a simple configuration of the MIT model of the North Atlantic Ocean can save up to 50% of the computational cost as compared to the optimization with full control space. The choice of the order-reduction has a significant impact on the convergence and the best reduction was obtained based on an Empirical Orthogonal Functions (EOF) analysis of a previous optimization using the same model configuration.

1 Introduction

Data assimilation aims to combine numerical models and observations to obtain the best description of the time-evolving state of a dynamical system. This technique, which is widely used in meteorology, has recently been pioneered in oceanography after the advent of oceanographic satellite missions and through the continuous progress of computational technology. The theoretical framework of data assimilation is well established (e.g., see [Wunsch(1996)]) and two main directions are usually followed in oceanography applications, one being a filter approach (e.g., Kalman filter), the other one being a smoother method (e.g., Lagrange multiplier or adjoint method). Filtering methods proceed by incrementally correcting the discrepancy between observations and a models prediction based on prior information about uncertainties in the model and the data. To the contrary, smoothing approaches seek to minimize the misfit between data and model trajectory in a "whole domain" approach, i.e., using data from the past and the future. To this end, the model-data misfit is reduced through adjustments of a well chosen set of control parameters.

Although smoother methods (RTF smoother and adjoint models) have been found to be very effective for complex data assimilation in oceanic GCMs, their computational burden remains one of the main obstacle to use them in a fully eddy-resolving setting or to transition existing state-of-the-art technology to operational centers. Especially the huge dimension of the system control vector (which in ongoing efforts on global scale over 10+ years has the order of 10^8 elements) makes the convergence of the adjoint optimization procedures very slow. One approach to reduce the cost of the adjoint method is therefore to reduce the dimension of the control vector. This may improve the convergence rate of the optimization since the later will be carried out then only in a subspace of much smaller dimension. In the context of an adjoint model this would also enable the use of a "full" Hessian matrix[2] in the optimization algorithm and therefore allow for the improvements in its performance.

Attempts related to reducing the computational burden of these algorithms have re-

[2] A diagonal approximation of the Hessian matrix is often used in Lagrange-multiplier approaches because of the large dimension of the system.

ceived considerable attention. As an example, different variants of the Kalman filter have been proposed to reduce the dimension of the system. The resulting "sub-optimal" filters essentially consist of the projection of the system state onto a low dimensional subspace using an order reduction operator [Cane et al.(1996), Cohn and Tolding(1996), Fukumori and Malanotte-Rizzoli(1995), Pham et al.(1997)]. An alternative direction that makes use of Monte Carlo methods to represent the error statistics given in the Kalman filter has also been developed [Evensen(1994), Burgers et al.(1998)]. In the context of the adjoint method, a simplification is been achieved by using a coarser grid in the adjoint model [Courtier et al.(1994)] or through truncation to reduce the integration cost of the adjoint model. As an example, [Vogeler and Schröter(1995)] filtered the adjoint gradients, only gradually allowing the cut off wavelength to become smaller and smaller. Yet another approach consists of performing searches for the optimum solution in the data space rather than the control space using the dual formulation (also called the Physical-space Statistical Analysis System (PSAS)[3]) of the variational methods [Courtier(1997)]. This approach, is equivalent to the Lagrange multiplier approach in the case of a linear model. It was shown to be very efficient in its convergence rate, since there is usually a large part of state space about which the observations provide no information at all [Bennett(1992), Wunsch(1996)].

Given the large – and ever increasing – number of observations in the current data assimilation systems, however, the implementation of these PSAS techniques require further simplification. This was achieved by introducing a reduced-order control space. [Vidard et al.(2000)] applied this approach to improve the convergence rate of the minimization while controlling the model error in the subspace spanned by the directions of the fastest growing perturbations. They also found that an order-reduction can be also beneficial for preventing the model from fitting into model noise and observation errors. Recently, [Durbiano et al.(2002)] showed the efficiency of this method in speeding up the optimization procedure while controlling the model initial conditions in a subspace generated by a set of Empirical Orthogonal Functions (EOFs) using a twin experiments

[3]The PSAS algorithm is equivalent to the representer method [Bennett(1992)].

approach.

A reduction of the control system dimension seems feasible in oceanic and atmospheric models because only a few modes are needed to represent most of the observed variability of these systems [De Mey(1997)]. Several methods to construct a reduced space are available in practice. For example one may apply the optimization on a coarser grid. Such a method, used by [Courtier et al.(1994)] and [Fukumori and Malanotte-Rizzoli(1995)], would however only reduce the dimension of the system by a factor of 10 or 100. Another possibility proposed by [Menemenlis and Wunsch(1997)] is to apply a series of transformations based on temporal, vertical and horizontal filters. Fast Fourier and Wavelet transforms are also alternatives to construct a low dimension approximation of the control vector.

In this paper, a statistical approach has been adopted, that of an Empirical Orthogonal Functions (EOF) analysis, which is widely used in oceanic and meteorological sciences. The EOF analysis essentially consists of a recursive construction of vectors, called EOFs, that maximize the correlation between the vectors. The reader may consult [Preisendorfer(1988)] for an exhaustive discussion of this approach. It has been used by [Cane et al.(1996)] in the context of reduced-order Kalman filtering and by [Durbiano et al.(2002)] in the context of reduced-order adjoint method.

The purpose of this paper is to test order-reduction schemes in a realistic setting which often proves to be much more complex and less forgiving than simple test of methods. Our goal is to obtain an optimization approach that would speed up the convergence of ongoing global synthesis approaches which are being performed as part of the "Estimating the Circulation and Climate of the Ocean " (ECCO) assimilation efforts (see [Stammer et al.(2002b)] for details). As part of the ECCO activities most available data have been assimilated into a global model using the MIT model with a 1° spatial resolution during the period 1992 through 2002. Control variables are typically initial temperature and salinity fields and daily fields of surface momentum, heat and freshwater fluxes, typically leading to 10^8 control variables [Köhl et al.(2004)].

In this present study we investigate the effect of an order-reduction scheme of the control vector in a smaller region of the sub-tropical Atlantic. Except for the domain, we used the same configuration as in [Stammer et al.(2002b)]. This work is a test of the convergence rate of the simplified approach which, if successful, will later on be implemented in the global efforts. We use EOFs to reduce the dimension of the atmospheric forcing fields which are used as control variables to adjust the model while assimilating real SST and SSH observations. To compute a "good" set of EOFs, the analysis has been applied on a priori optimized set of control vectors.

As will become clear below, a reduced-order scheme shows a faster initial convergence, but would fail to reduce the model-data misfit as much as is achieved using the full control space. This is because the reduced space does miss some important smaller scale features that are not represented in the available prior statistics. A new strategy is therefore proposed and tested here, which consists of starting the optimization in the reduced space to speed up the convergence at the early optimization steps and than to continue the optimization in the full control space to capture the missing scales in the reduced space.

Section 2 briefly summarizes the formulation of the adjoint method and describes the order reduction strategy. The design and the results of the numerical experiments are presented in Sections 3. Finally, concluding remarks are given in Section 4.

2 The Approach

The Lagrange-multiplier (adjoint) assimilation method is a particular case of the more general framework of the optimal control theory and belongs to the so called "variational" or "whole domain" approaches. Those methods consist of finding the model trajectory which best fits available data over a given period of time. The model equations are added as constraints to the optimization problem. This leads to an optimization problem of a cost function, generally denoted by J, measuring the discrepancy between the model solution and data with regard to a well-chosen set of control variables. The gradient of the cost function with respect to control variables, which are provided through the adjoint model, are used to determine descent directions toward the minimum in an iterative procedure. [Wunsch(1996)] provides all details on the theory and applications in oceanographic problems.

To reduce the order of the control vector we follow [Fukumori and Malanotte-Rizzoli(1995)], denoting by $\mathbf{B}$ the order reduction operator that projects the control vector onto some truncated basis set

$$\tilde{\mathbf{u}} = \mathbf{B}\mathbf{u}, \tag{1}$$

where $\tilde{\mathbf{u}}$ is a linear approximation of the original control vector $\mathbf{u}$ with a smaller dimension. A pseudo-inverse operator $\mathbf{B}^*$, which can be thought as a reconstruction operator that maps the reduced control vector on the original control space, is also defined such that

$$\mathbf{B}\mathbf{B}^* = \mathbf{I}, \tag{2}$$

where $\mathbf{I}$ is the identity matrix. The variational assimilation problem can then be simplified to a so-called reduced Lagrange-multiplier that consists of seeking for the reduced control vector $\tilde{\mathbf{u}}$ solution of the variational optimization problem. The gradient of the cost function J with respect to $\tilde{\mathbf{u}}$ can be obtained from the "original" adjoint model by projecting the "full" gradient onto the reduced control space according to

$$\left(\frac{\partial J}{\partial \tilde{\mathbf{u}}}\right)^T = \mathbf{B}\left(\frac{\partial J}{\partial \mathbf{u}}\right)^T. \tag{3}$$

The appeal of this order-reduction approach is that no changes are required for the model and its adjoint, only the optimization is now applied in a reduced space. The original adjoint model is therefore run forward and backward with the original resolution and the order-reduction algorithm is applied to the resulting adjoint gradients prior and after each linear search step. We note that any computational savings will therefore come in this approach only through the saving in number of iterations required to obtain a converged solution.

From a practical point of view, the advantage is that only two simple steps have to be added to the entire optimization procedure and that in a very simple manner one can go from a full order to a reduced order and vice versa: in a first step the solution of the adjoint model is being projected onto the reduced control space using (3); in the second step the full order control vector is being reconstruct (computed in the reduced space with the new iteration of the optimization algorithm) into the original control space using the approximated formula

$$\mathbf{u} \approx \mathbf{B}^{*}\tilde{\mathbf{u}}. \tag{4}$$

It is worthwhile to note that the cost of these two operations is negligible with regard to the integration cost of the numerical model (forward and backward).

To achieve an order reduction of the control space, EOFs were calculated from a set of realizations of time series of the forcing fields. The EOFs depend therefore on space and time and contain statistical information about the spatial and temporal coherence of the control variables, which could enhance the performance of the optimization procedure. The EOFs, however, can only extract the observed variations of the time series they represent. The set of vectors from which the EOFs will be computed should be therefore representative of the variability of the variables used to control the model.

The choice of the set of vectors from which the EOFs are computed is a key issue for building an efficient reduced adjoint method, since the performance of this method highly depends on the representativeness of the reduced control space. Following a twin experiments approach, [Durbiano et al.(2002)] used a long time series of model outputs to determine a reduced space for the initial conditions. Although the results of this approach were found to be very effective in twin experiments, the resulting EOFs may be not as effective when using real data because the variability of the corrections imposed by the real observations is often substantially different from that simulated by the model. In our case, where the atmospheric forcing are considered as control variables, the EOFs can be either computed from a time series of forcing fields (e.g. NCEP) over many years or from a priori optimized set of control vectors. The former assumes that the forcing corrections are part of the space of the forcing variability whereas the later seeks the corrections in the search space of a prior optimization, relying on the assumption that the search paths of different optimization largely share a common space.

Two reduced-order control spaces were first determined by applying a multivariate EOF analysis on: (TS1) 10 years of NCEP forcing fields (which are used to force the model), and (TS2) 10 years of available ECCO adjustments to the NCEP forcing fields

[Stammer et al.(2002b)]. The later have been optimized by a global state estimation procedure calculated by [Stammer et al.(2002b)] on a $2^\circ \times 2^\circ$ global grid using the MIT model with a similar configuration to the one used in the current experiments. These two sets obviously provide two samples of 10 (one from each year) vectors to compute the EOFs. The EOFs computed from the ECCO solution are expected to be more appropriate for our problem since these adjustments have been already optimized using the classical adjoint method. A third set of EOFs was computed from a set of vector (TS3) taken from priori iterations with our model. More precisely, to determine TS3, the model has been adjusted over a period of one year starting from January 1992 with the classical adjoint method and the first 40 iterations resulting from this experiment were saved. These three sets of EOFs will be used to study to which extend a good EOFs reduced space is related to the choice of the model and its domain.

3 Experiments

We used the ECCO ocean general circulation model, which is derived from the MIT model ([Marshall et al.(1997)]). An adjoint code to the forward model was obtained from the automatic differentiation tool of [Giering and Kaminski(1998)]; see also, [Marotzke et al.(1999)]. To test the reduced-order adjoint assimilation approach, we investigate a simple box of the Sub-tropical North Atlantic ocean extending from $10^\circ N$ to $40^\circ N$ and from $42^\circ W$ to $4^\circ E$. Zero fluxes of heat and salt and free-slip conditions are applied at the lateral boundaries, which are closed. Non-slip conditions are applied at the bottom. The model is set up on a $2^\circ \times 2^\circ$ horizontal grid and 23 vertical levels, with the first 6 levels in the upper 100 meters. The model domain has a realistic bottom topography based on the [ETOPO5(1988)] dataset. Free-slip bottom boundary conditions and non-slip boundary conditions at lateral walls are used. Laplacian viscosity and diffusivities are imposed, with $\nu_h = 1 \times 10^4$ m^2/s and $\kappa_h = 10^2$ m^2/s and $\nu_v = 10^{-3}$ m^2/s and $\kappa_v = 10^{-5}$ m^2/s, in the horizontal and vertical, respectively. The mixed layer is modeled with the "KPP" code of [Large et al.(1994)]. The time step is 1 hour. Atmospheric forcing consists of daily heat and fresh water fluxes, and twice-daily zonal and meridional wind stress components from the National Center for Environmental Prediction (NCEP)/National Center for Atmospheric Research (NCAR) re-analysis project [Kalnayand et al.(1996)]. The model is started from rest and Levitus temperature T and salinity S fields. Because no restoring was applied at the closed boundaries, the model provides realistic simulation only for limited periods of time. See [Stammer et al.(2002b)] for a global setup.

The adjoint model has been constrained with daily along track sea surface height (SSH) from TOPEX/POSEIDON and ERS data, monthly Reynolds surface temperature (SST) and subsurface Levitus S and T data. The control vector consists of heat flux, fresh water flux, and the wind stress which were adjusted every two days. Data and model errors are prescribed only along the diagonal of the error covariances and are the same as in [Stammer et al.(2002b)]. They have been approximated by the error profiles for temperature and salinity taken from Levitus data and by 50% of the data variability for the SSH. Prior errors are provided as standard deviation (STD) of the differences between NCEP and QuickSCAT scatterometer wind fields for the wind stress and as 1/3 of the local STD of the NCEP forcing for the net heat and freshwater fluxes. Note that the inclusion of the model initial conditions in the control vector is strait forward and do not require any special treatment. It has not been adjusted in this study in order to be flexible and allow for other reduction approaches that are not discussed here. The descent directions toward the minimum have been determined using Quasi-Newton M1QN3 developed by [Gilbert and Le Maréchal (1989)]. The first guess for the optimization are provided by the NCEP forcing fields which therefore were used as the background fields in the cost function, we estimate deviations from these fields which we refer to as corrections.

After a first experiment numerical conducted in 1992 over a one year period to determine a priori set of control vectors for the computation of EOFs, experiments have been carried out subsequently over a period of one year starting from January 1993 to evaluate the performance of the reduced order optimization strategy and to study its sensitivity to the choice of EOFs.

3.1 Sensitivity to the choice of EOFs

The reduced adjoint method has been first implemented with the three sets of EOFs computed from (TS1) the NCEP forcing and (TS2) the ECCO adjustments, and (TS3) priori estimates in 1992, as described in section 2. The goal of this comparison is to study the performance of this method with regard to the choice of the reduced control space. 4 EOFs were retained from each EOF analysis, they almost describe 69%, 73% and 90% of the total variance, respectively. The cost-functions for these experiments are plotted in Fig. 1. It can be seen that the performance of the reduced adjoint method highly depends on the reduced space. Indeed, the cost-function decrease stagnates after only 8 iterations when using EOFs from NCEP forcing, suggesting that a large fraction of the variability of the forcing adjoint corrections was not well represented by the variability of the NCEP forcing fields. The use of the EOFs computed from the ECCO adjustments clearly improves the assimilation results of the reduced adjoint method. Surprisingly, the overall performance of this method remains however unsatisfactory. The best assimilation results were clearly obtained when the reduced adjoint method was implemented with EOFs computed from the a priori optimized control vectors in 1992. Such a reduced space improves the cost function decreasing by more than 30% comparing to the other set of EOFs. This result can be expected since these 'analyzed' vectors contain information from both the model and the observations and are adjusted on the same domain, definitely providing the best available set to determine the reduced control space.

3.2 Sensitivity to the number of retained EOFs

As demonstrated, a massive order reduction guarantees a fast convergence rate of the optimization algorithm. This would however tend to miss some important features (generally related to small scale phenomena) even if the "representativeness" error of the control vector in the reduced space is very small. The efficiency of the reduced adjoint method is therefore limited by the "amount" of variability represented by the reduced space. As confirmed in our numerical experiments, the decrease of the cost function in the reduced space stagnates quickly after all information contained in the reduced space has been used. The reduction of the control vector dimension obviously comes at the cost of introducing an approximation.

A second series of experiments was therefore conducted to study the sensitivity of the convergence of the reduced adjoint method to the number of retained EOFs (dimension of the reduced space). Fig. 2 shows the evolution of the total cost function as a function of the number of iterations obtained with the reduced adjoint method using 2 EOFs, 4 EOFs and 6 EOFs (computed from the 1992 iterations (TS3)), which almost describe 76%, 90% and 92% respectively. The optimization was therefore performed in subspaces of dimension 2, 4 and 6 instead of 284700. As expected, the cost function decreases rapidly during the first few iterations as the dimension of the reduced space decreases. The overall performance (cost-function) of the reduced adjoint method is however improved when more EOFs are used. Therefore the number of retained EOFs should be chosen cautiously: small enough to make sure of fast convergence rate and large enough for an efficient representation of the variability of the full control space. A comparison with Fig. 4 reveals that the performance with the 6 EOFs is still not satisfactory as compared to the solution of the full control space, suggestion that some hybrid version of both is required as final approach.

3.3 Performance of the reduced adjoint method

A simple way to improve the performance of the reduced adjoint method is to continue the optimization in the full control space for a few more steps in order to capture the missing variability in the reduced space. Starting the optimization in the reduced space would definitely speed up the convergence rate at the early minimization process. Furthermore, variational method search for the solution iteratively, starting by determining the components which are the richest in terms of information (often related to large-scale structures). Then the components with less information is determined. This strategy can be therefore thought of as a way to get the solution of the classical adjoint method but with fewer iterations as shown in Fig. 3.

To study the performance of the new strategy, three experiments were conducted with: (E1) the classical adjoint, (E2) the reduced adjoint method, and (E3) the classical adjoint for the first 3 optimization iterations and then with the reduced adjoint method. In the last two experiments, the reduced space is generated by the EOFs computed from the 1992's iterations (TS3) as described in section 2. The solution obtained with the classical adjoint method is used as a reference to evaluate the performance of the reduced optimization strategy. In the following 4 EOFs were retained for (TS3) describing 90% of the total variance of the priori estimates in 1992.

Fig. 4 shows the evolution of the total cost function as function of the number of iterations for the three experiments. It can be seen that the classical adjoint method behaves fairly well, decreasing the cost function by a factor of 4. The convergence of the optimization procedure seems also to be reached after only 40 iterations. While the cost function from (E2) decreases rapidly during the first 10 iterations, the solution obtained by this strategy is not as good as the one provided by the classical adjoint method, suggesting that some features (often related to small scales) were not represented by the reduced control space (generated by the EOFs). This also suggest that the performance of the optimization scheme in the reduced control space can be, and will always be in practice, limited by the representativeness of this sub-space. The results of (E3) show that by pursuing the optimization in the full space (after 3 iterations in the reduced space) one can get a quite similar solution to the one obtained in (E1) but with less iterations. For instance, the solution obtained in (E1) after 20 iterations can be obtained in (E3) after only 10 iterations. This can be also seen from Fig. 5, which plots the mean adjustment to the heat flux and the zonal component of the wind stress after 20 iterations with (E1) and compare it to the solution obtained after 10 iterations with (E1) and 10 iterations with (E3). Indeed, the corrections that result after 20 iterations with (E1) and 10 iterations with (E3) are similar and the difference seems to be only governed by small scales. The adjustments obtained after 10 iterations with (E1) are still quite different and require more iterations to reach a satisfactory solution.

Faster convergence of the cost function was also obtained by (E3) for the individual misfit terms of the cost function. The decreasing is however shown to be different for each variable (Fig. 6). While the convergence to the Reynolds and Levitus temperature (T)

data seems to be reached after only 20 iterations with the classical adjoint method, more iterations are needed for the salinity (S) and the particularly for the sea-surface height (SSH). This is due to the large weights used for the T data making the corresponding misfit terms dominant in the total cost function with regard to the other terms. The improvement of the reduced optimization strategy to the convergence to Levitus salinity data at the first iterations can be also expected since the optimization has only 4 degrees of freedom. The slow convergence to the SSH in the reduced space suggests a bad representativeness of the descent directions of this variable in the EOFs. Indeed, the good representation of the descent directions for the 1993 S and T measurements in the EOFs computed from 1992 corrections is probably due to the climatological character of these data. Finally, similar conclusions can be made from Fig. 7. As it can be seen, the averages of the sea-surface temperature (SST) over the assimilation period from the former runs are in good agreement with the Reynolds data. The solution of the classical adjoint method is also shown to be quite well estimated after only 10 iterations. The solution from the classical adjoint method after 20 iterations and from the reduced-order adjoint method after 10 iterations are very similar. The propagation of the observations to the sub-surface layers is also shown to be faster with the reduced adjoint method (Fig. 7). The adjoint solutions after 10 iterations in the full and the reduced space are shown to have similar patterns but with different magnitude, suggesting similar descent directions in both cases at the first iterations but with faster convergence of the reduced adjoint method.

4 Discussion

We investigated the potential of an order reduction scheme to reduce the cost of four-dimensional data assimilation methods by applying the minimization in a reduced control space rather than the full control space. This is important because computational and memory burden remain the main obstacle for the use of adjoint assimilation methods for data assimilation in realistic high resolution oceanic models and in operational settings. The control vector is composed in this example of only the surface forcing fields. However, results can easily be extended to include also the models initial temperature and salinity conditions. Although the assimilation experiments were conducted here in a simple configuration of the North Atlantic ocean, they do demonstrated the effectiveness of this optimization strategy.

The following findings can be summarized:

- Using a reduced-order control space significant speed up of the convergence can be found. In this study, the control vector is composed of the surface forcing fields. The choice of EOFs of the control vector was essential in determining the convergence. The best convergence was found using EOFs from a previous optimization performed over the same model domain, this approximating in a best possible way the graves modes of the surface forcing uncertainties and uncertainties of the specific model configuration used during the study.

- Using a reduced-order control space does speed up the initial convergence, but never reaches the quality of the full optimization. A hybrid approach is therefore used which starts with a reduced-order control space. The initial iterations do improve the model-data misfit significantly based on the gravest modes of the control space. The full control space is used subsequently to adjust also smaller scales required for the assimilation of real data in the presence of boundary currents or smaller time scale fluctuations. The new strategy can be therefore thought of as a way to get the solution of the classical adjoint method but with fewer iterations (as described in Fig. 3).

- Using the hybrid approach, roughly 50% of the iterations could be saved in our simple setting and results are very close to those obtained using the full control space. The reduced optimization strategy was found to be fairly effective in assimilating SST and SSH data while controlling forcing fields, leading to a similar solution to the one obtained with the classical optimization procedure but half the number of iterations.

The most important factor determining the performance of the proposed method is the choice of the reduced space. The space has been computed by applying an EOF analysis on a series of forcing fields from the NCEP reanalysis and a priori optimized set of control vectors, both on a global scale and in the regional model domain. A significantly better performance was found for the latter case which accounts for both forcing errors and errors of the regional model set up alike and provides a significant speedup in comparison to the classic approach. The drawback of this choice is that a previous assimilation experiment is needed. Though this seems seldom the case for most state estimation experiments, it is the general case for an operational applications in which the period of assimilation keeps being extended as new data comes in.

Acknowledgments

Re-analysis surface forcing fields from the National Center for Environmental Prediction/National Center for Atmospheric Research (NCEP/NCAR) are obtained through a computational grant at NCAR. The adjoint compiler TAF from the company FastOpt was used to create the adjoint general circulation model. The computational support through a NRAC grant from the National Partnership for Computational Infrastructure (NPACI) and through a grant from the National Center for Atmospheric Research is acknowledged. Supported in part through NASA grant NAG5-7857, through NSF grant OCE 9730071 and through two contracts with the Jet Propulsion Laboratory (958125 and 1205624). This is a contribution of the Consortium for Estimating the Circulation and Climate of the Ocean (ECCO) funded by the National Oceanographic Partnership Program under grants N00014-99-1-1049.

References

[Bennett(1992)] Bennett A. (1992): Inverse methods in physical oceanography. *Cambridge University Press*, 346 pp.

[Cane et al.(1996)] Cane M.A., A. Kaplan, R.N. Miller, B. Tang, E.C. Hackert and A.J. Busalacchi (1996): Mapping tropical Pacific sea level: data assimilation via a reduced state Kalman filter. *J. Geophys. Res.*, **101**, no.C10, 22599-22617.

[Burgers et al.(1998)] Burgers G., P.J. Van Leeuwen and G. Evensen (1998): Analysis scheme in the ensemble Kalman filter. *Mon. Wea. Rev.*, **126**, 1719-1724.

[Cohn and Tolding(1996)] Cohn S.E. and R. Tolding (1996): Approximate Kalman filters for stable and unstable dynamics. *J. Meteor. Soc. Japan*, **74**, 63-75.

[Courtier(1997)] Courtier P. (1997): Dual formulation of four-dimensional variational assimilation. *Quart. J. Roy. Meteor. Soc.*, **123**, 2449-2461.

[Courtier et al.(1994)] Courtier P., J.-N. Thèpaut, and A. Hollingsworth (1994): A strategy for operational implementation of 4D-Var, using an incremental approach. *Quart. J. Roy. Meteor. Soc.*, **120**, 1367-1387.

[De Mey(1997)] De Mey P. (1997): Data Assimilation at the oceanic mesoscale: A Review. *J. Met. Soc. Japan*, **75**(1B), 415-427.

[Durbiano et al.(2002)] Durbiano S., E. Blayo, J. Verron, J. Blum, F.-X. L. Dimet (2002): A reduced order strategy for 4D-Var ocean data assimilation. Submitted to *Quart. J. Roy. Meteorol. Soc.*

[Evensen(1994)] Evensen G. (1994): Sequential data assimilation with a nonlinear quasi-geostrophic model using Monte Carlo methods to forecast error statistics. *J. Geophys. Res.*, **99**, no.C5, 10143-10162.

[ETOPO5(1988)] ETOPO5, 1988: Digital relief of the surface of the earth. worldwide bathymetry/topography data announcement 88-mgg-02, *Technical report, National Geophysical Data Center*.

[Fukumori and Malanotte-Rizzoli(1995)] Fukumori I. and P. Malanotte-Rizzoli (1995): An approximate Kalman filter for ocean data assimilation: an example with an idealized Gulf Stream model. *J. Geophys. Res.*, **100**, 6777-6793.

[Ghil and Malanotte-Rizzoli(1991)] Ghil M. and P. Malanotte-Rizzoli (1991): Data assimilation in meteorology and oceanography. *Adv. Geophys.*, **33**, 141-266.

[Gilbert and Le Maréchal (1989)] Gilbert J. C. and C. Le Maréchal (1989): Some numerical experiments with variable storage Quasi-Newton algorithms. *Math. Program.*, **45**, 407-435.

[Giering and Kaminski(1998)] Giering R. and K. Thomas (1998): Recipes for adjoint code construction. *ACM Trans. On Math. Software*, **24**, 437-474.

[Kalnayand et al.(1996)] Kalnayand, E., M. Kanamitsu, R. Kistler, W. Collins, D. D. L. Gandin, M. Iredell, S. Saha, G. White, J. Woollen, Y. Zhu, M. Chelliah, W. Ebisuzaki, W. Higgins, J. Janowiak, K. C. Mo, C. Ropelewski, A. Leetmaa, R. Reynolds, and R. Jenne (1996): The NCEP/NCAR re-analysis project. *Bull. Amer. Meteor. Soc.*, **77**, 437-471.

[Köhl et al.(2004)] Köhl A., D. Stammer and B. Cornuelle (2004): Interannual to Decadal Changes in the ECCO Global WOCE Synthesis. *J. Phys. Oceanogr.* In preparation.

[Large et al.(1994)] Large W., J. McWilliams, and S. Doney (1994): Oceanic vertical mixing: A review and a model with a nonlocal boundary layer parameterization. *Reviews of Geophysics*, **32**, 363-403.

[Losch and Wunsch(2003)] Losch M. and C. Wunsch (2003): Bottom topography as a control variable in ocean model. *J. Atm. Ocean Tech.*, in press.

[Marotzke et al.(1999)] Marotzke, J. R. Giering, Q. K. Zhang, D. Stammer, C. N. Hill and T. Lee (1999): Construction of the adjoint MIT ocean general circulation model and application to Atlantic heat transport sensitivity. *J. Geophys. Res.*, **104**, 29,529-29,548.

[Marshall et al.(1997)] Marshall J., A. Adcroft, C. Hill, L. Perelman and C. Heisey (1997): A finite-volume, incompressible Navier Stokes model for studies of the ocean on parallel computers. *J. Geophys. Res.*, **102**, 5753–5766.

[Menemenlis and Wunsch(1997)] Menemenlis N. and C. Wunsch (1997): Linearization of an oceanic general circulation model for data assimilation and climate studies. *J. Atmos. Ocean Tech.*m **14**, 1420-1443, 1997.

[Pham et al.(1997)] Pham D.T., J. Verron and M.C. Roubaud (1997): Singular evolutive Kalman filter with EOF initialization for data assimilation in oceanography. *J. Mar. Syst.*, **16**, 323-340.

[Preisendorfer(1988)] Preisendorfer R. (1988): Principal component analysis in meteorology and oceanography. *Elsevier Sci. Publ.*, **17**, 425 pp.

[Stammer et al.(2002b)] Stammer, D., C. Wunsch, R. Giering, C. Eckert, P. Heimbach, J. Marotzke, A. Adcroft, C.N. Hill, and J. Marshall, The global ocean circulation during 1992 –1997, estimated from ocean observations and a general circulation model, *J. Geophys. Res.*, 107(C9), 3118, doi:10.1029/2001JC000888, 2002.

[Thacker and Long(1988)] Thacker W. and R. Long (1998): Fitting dynamics to data. *J. Geophys. Res.*, **93(C2)**, 1227-1240.

[Verlaan and Heemink(2001)] Verlaan M. and Heemink A.W. (2001): Nonlinearity in data assimilation applications: A practical method for analysis. *Mon. Wea. Rev.*, **129**, 1578-1589.

[Vidard et al.(2000)] Vidard P.A., E. Blayo and F.X. Le Dimet (2000): 4D variational data analysis with imperfect model. *Flow. Turbul. Combust.*, **65**, 489-504.

[Vogeler and Schröter(1995)] Vogeler and Schroter (1995): Assimilation of satellite altimeter data into an open ocean model. *J. Geophys. Res.*, **100**, 15951-15963.

[Vukicevic and Bao(1998)] Vukicevic T. and Bao J.W. (1998): The effect of linearization errors on 4D-Var data assimilation. *Mon. Wea. Rev.*, **126**, 1695-1706.

[Wunsch(1996)] Wunsch C. (1996): The ocean circulation inverse problem. *Cambridge Univ. Press*, New York, 442pp.

[Yang et al.(1996)] Yang W.Y., I.M. Navon and P. Courtier (1996): A new Hessian preconditioning method applied to variational data assimilation experiments using NASA general circulation models. *Mon Wea. Rev.*, **124**, 1000-1017.

List of Figures

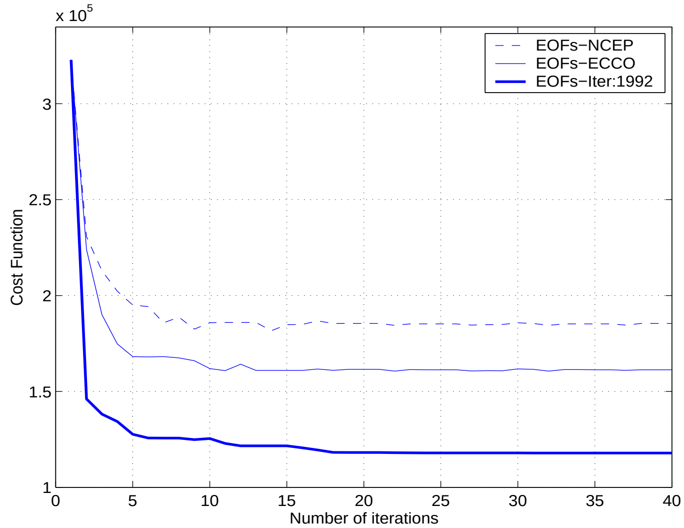

Figure 1: Evolution of the total cost function as function of the number of iterations obtained with the reduced adjoint method using 4 EOFs computed from NCEP forcings (TS1), ECCO forcings (TS2) and priori iterations in 1992 (TS3).

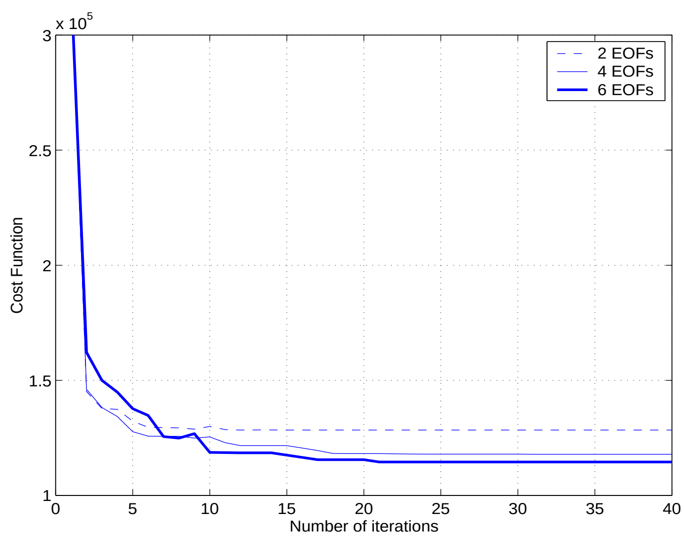

Figure 2: Evolution of the total cost function as function of the number of iterations obtained with the reduced adjoint method using 2 EOFs, 4 EOFs and 6 EOFs.

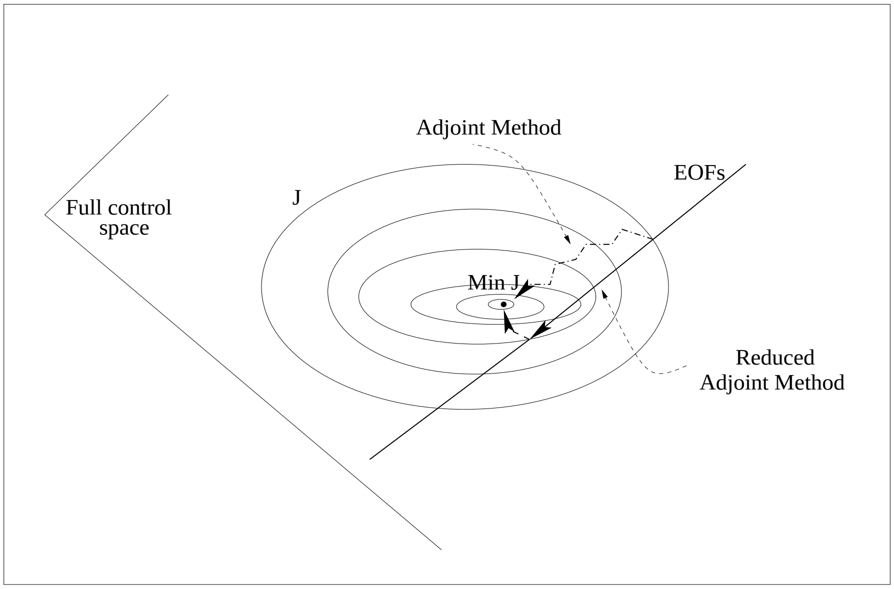

Figure 3: Schematic of the decreasing of the cost function in the full and reduced space.

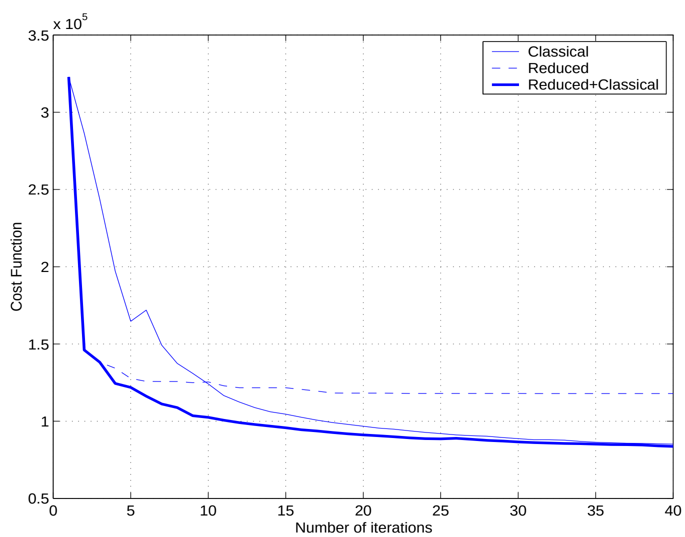

Figure 4: Evolution of the total cost function as function of the number of iterations from the (i) classical adjoint method, (ii) the reduced adjoint method (using 4 EOFs), and (iii) the reduced adjoint method (using 4 EOFs) at the first 3 optimization iterations and then the classical adjoint method.

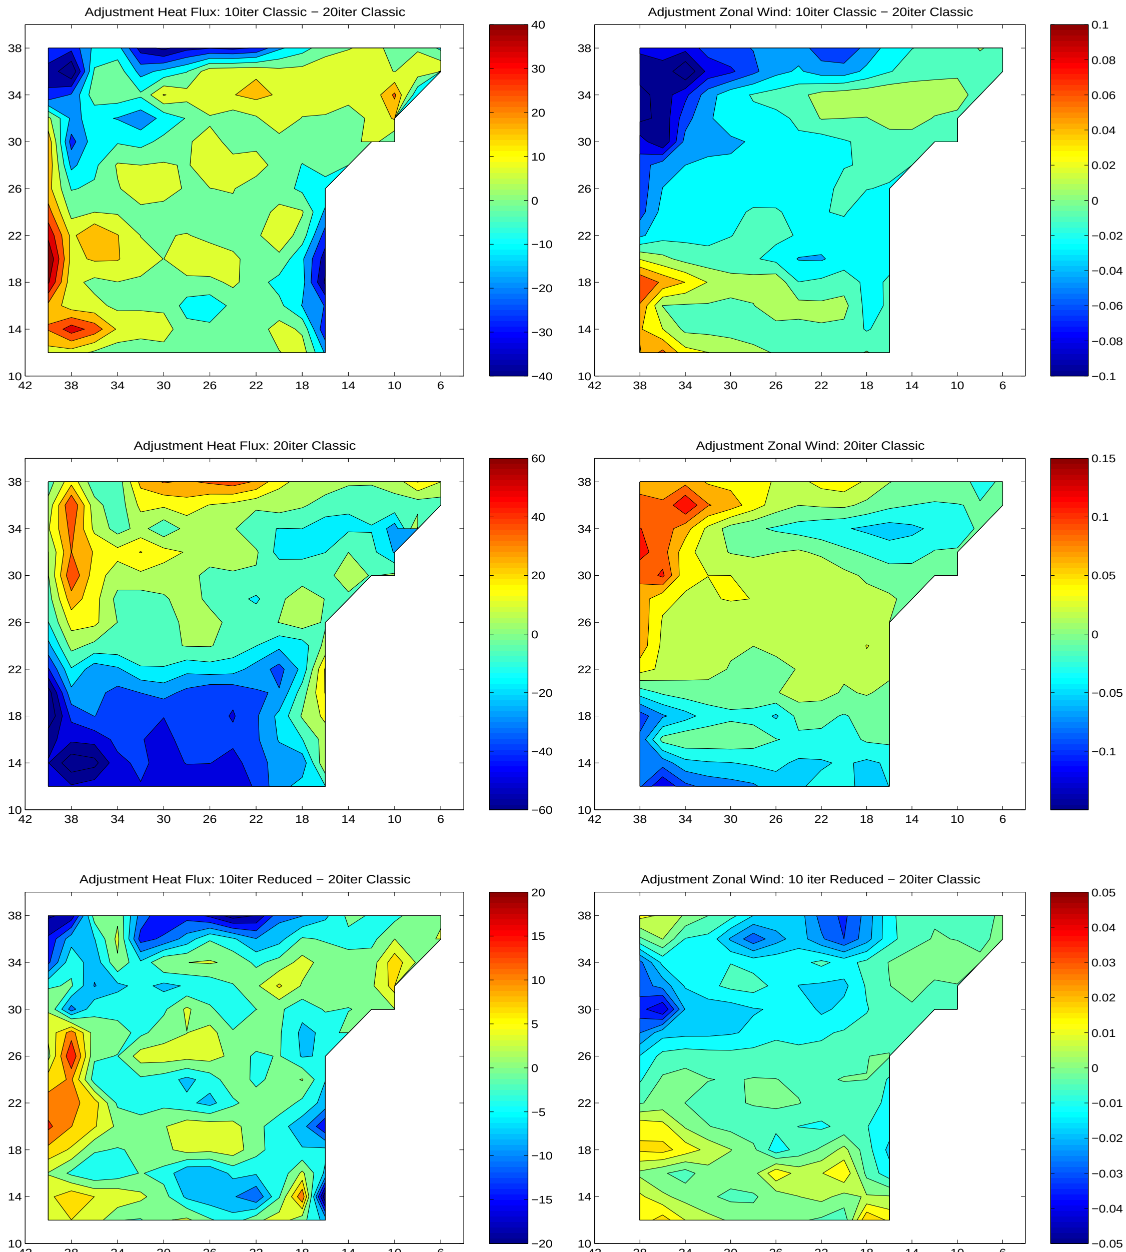

Figure 5: Difference of the time average (over the one year assimilation period) between the adjustments to the heat flux (left panel) and the zonal component of the wind stress (right panel) after 20 iterations and 10 iterations with the classical adjoint method (top), and 20 iterations with the classical adjoint method and 10 iterations with the reduced adjoint method (bottom). The middle figures show the time average of the adjustments to the heat flux (left panel) and the zonal component of the wind stress (right panel).

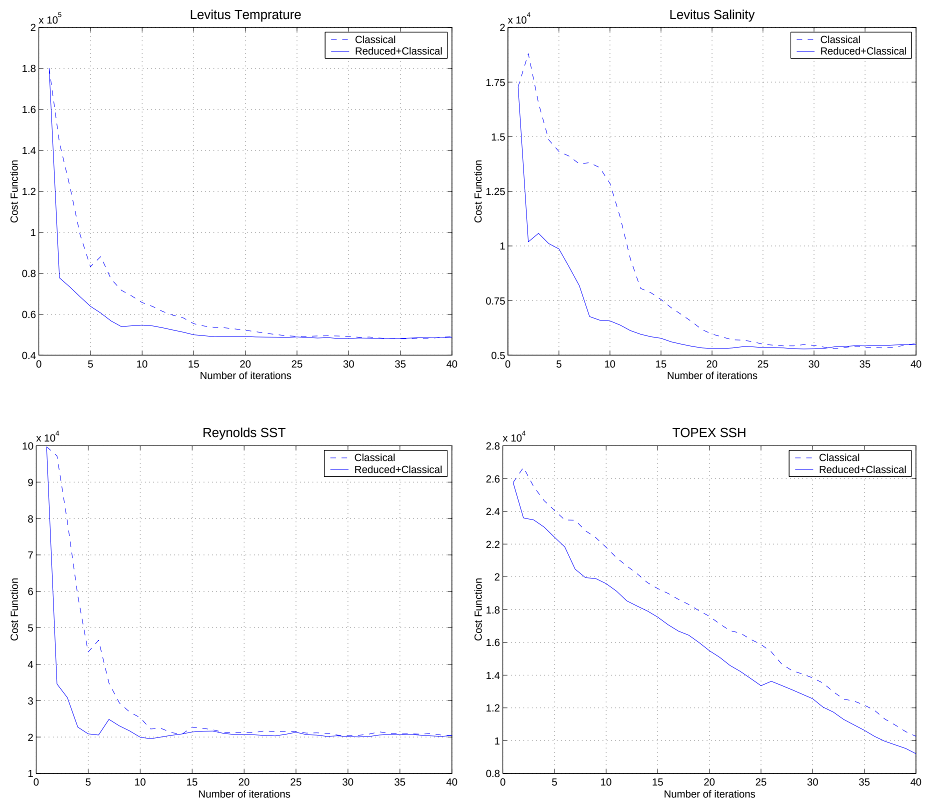

Figure 6: Evolution of the individual misfit terms of the cost function as function of the number of iterations from the classical adjoint method and the reduced adjoint method (using 4 EOFs) at the first 3 optimization iterations and then the classical adjoint method.

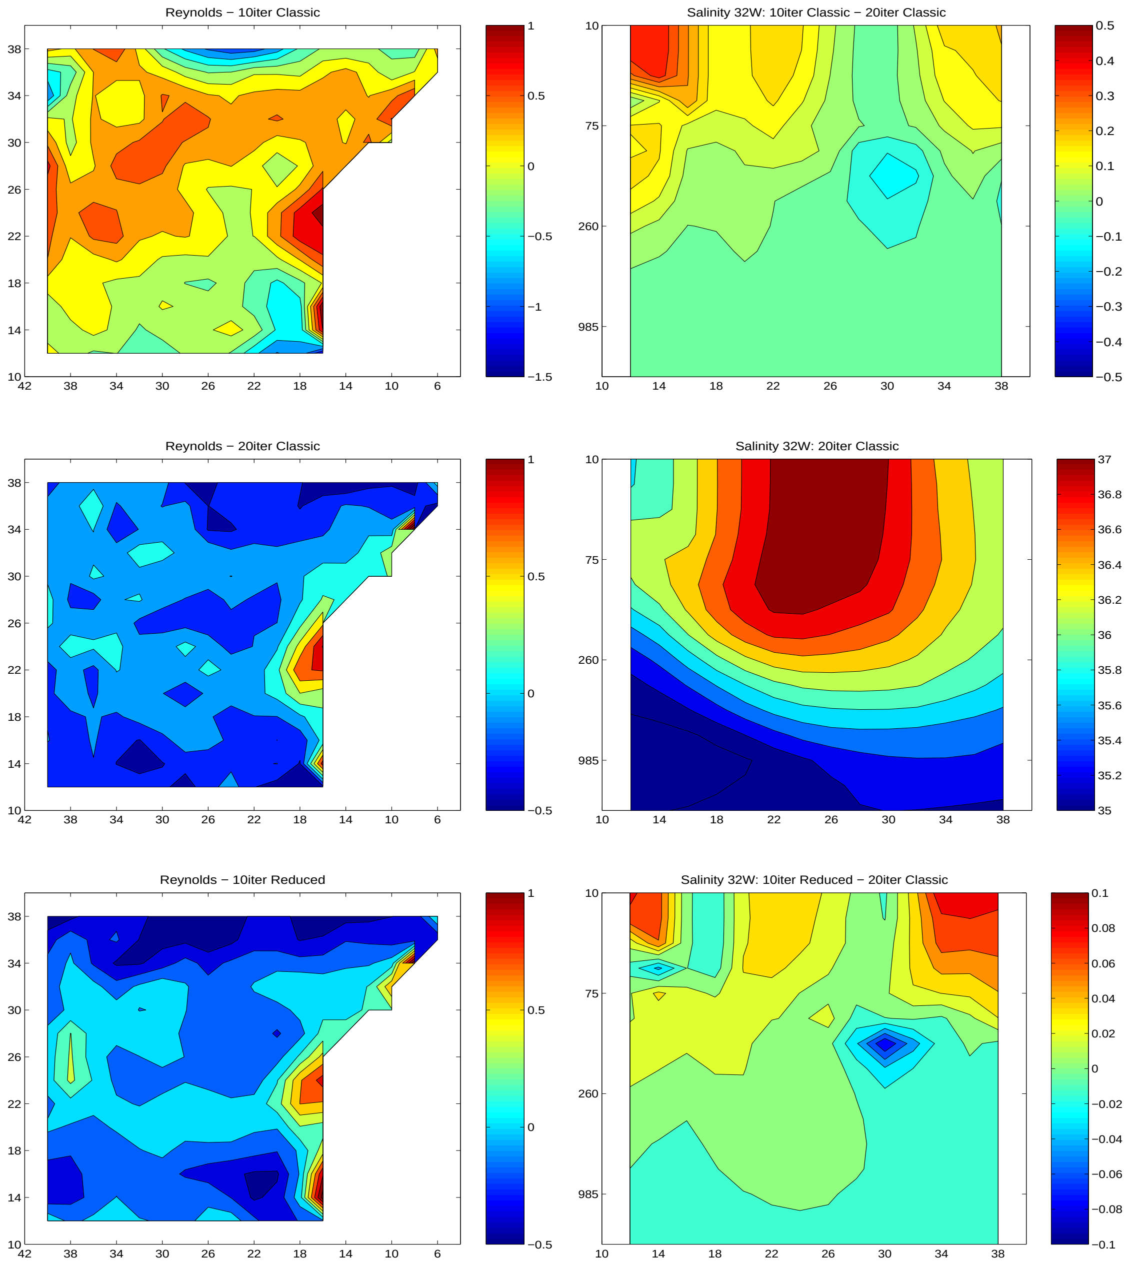

Figure 7: *Right panel*: Difference of the time average between Reynolds data and estimated SST after 10 iterations (top), 20 iterations (middle) with the classical adjoint method, and 10 iterations with the reduced adjoint method (bottom). *Left panel*: Difference of the time average between salinity cross-section at 32°W after 20 and 10 iterations with the classical adjoint method (Top), and 20 iterations with the classical adjoint method and 10 iterations with the reduced adjoint method (bottom). The middle figure shows the average of the salinity cross-section at 32°W after 20 with the classical adjoint method.