JOURNAL OF OPERATIONAL OCEANOGRAPHY, 2017
http://dx.doi.org/10.1080/1755876X.2017.1306376

Taylor & Francis
Taylor & Francis Group

A comparison of two preconditioner algorithms within the representer-based four-dimensional variational data assimilation system for the Navy coastal ocean model

Innocent Souopgui[a], Hans Ngodock[b], Matthew Carrier[b] and Scott Smith[b]

[a]Department of Marine Sciences, The University of Southern Mississippi, Stennis Space Center, MS, USA; [b]Naval Research Laboratory, Code 7320, Stennis Space Center, MS, USA

ABSTRACT

Four-dimensional variational (4D-Var) data assimilation for operational systems requires the solution of large linear systems that are poorly conditioned in general. In addition to efficient iterative solvers for linear systems, using a good preconditioner is required to guarantee an acceptable solution with a small number of iterations. We consider the assimilation of ocean observations using a weak constraints 4D-Var for the Navy Coastal Ocean Model (NCOM), based on the representer method. Two methods of preconditioning the linear system are implemented, namely the scaling of the linear system by the square root of the observations error variances, and the approximating stabilised representer matrix based on the computation of some representer functions. We evaluate their convergence using criteria such as the norm of the residuals and of the gradient of the cost function, the analysis error evaluated at the observation locations, and finally the convergence of the sequence of analyses. Results from all criteria show that the fastest convergence is achieved with the rescaling preconditioner when the conjugate gradient used to solve the linear system is equipped with a suitable inner product instead of the Euclidean.

ARTICLE HISTORY
Received 1 June 2015
Accepted 9 March 2017

1. Introduction

The four-dimensional variational (4D-Var) data assimilation algorithm seeks to minimise a cost function defined as a weighted sum of squared discrepancies between the model and the observations, using the model dynamics as a constraint, either strong or weak. Applications of 4D-Var to geophysical models of the atmosphere or the ocean are usually computationally expensive, primarily due to the necessity to solve the adjoint and the nonlinear/linearised models multiple times, and the size of the model state vector. The latter arises from a large model domain, or high horizontal and vertical resolutions, or both. The iterative process by which the 4D-Var algorithm minimises the cost function is usually poorly conditioned. Operational implementation of the 4D-Var necessitates that an accurate analysis be reached within a limited wall clock time constraint, that is, with the fewest iterations possible. This requires, at least, a monotonic decrease at every iteration of the cost function itself or the residuals of the linear system associated with the minimisation of the cost function. Given the fact that each iteration requires the integration of the adjoint and either the forward nonlinear or tangent linear models, the use of a preconditioner becomes imperative in order to mitigate the cost of the minimisation process and render the 4D-Var computationally affordable for operational use. This paper investigates the implementation of existing preconditioners and compares their relative performances within a weak constraints 4D-Var system using a realistic ocean model.

The minimisation of the cost function is usually carried out either in the control space or in the observations space. In the control space (usually taken to be the space of initial conditions in the classic 4D-Var), the iterative process through which the minimum of the cost function is obtained can be preconditioned by the square root of the background error covariance, Lorenc (1988), see also Courtier (1997), or by strategies such as the one proposed by Robert et al. (2006), that consists of first solving the 4D-Var problem in a reduced state space and using the solution as the initial guess for the full problem.

Here we consider the minimisation of the cost function in the observations space since it is usually much smaller than the control space, especially when a weak constraints 4D-Var is adopted. Thanks to the properties

CONTACT Hans Ngodock hans.ngodock@nrlssc.navy.mil Naval Research Laboratory, Code 7320, Stennis Space Center, MS, USA

(symmetric and positive definite) of the matrix of the linear system (1), the conjugate gradient (CG) algorithm of Hestenes and Stiefel (1952) is a good candidate for a solver, even though linear such systems are usually ill conditioned in the context of data assimilation. The literature on building preconditioners for linear systems is abundant, see for example Golub and Van Loan (2013, chapter 11) and Saad (2003, chapter 10). However it mostly applies to linear systems that have explicit matrices, which is not the case in data assimilation where the computation and storage of the matrix of the linear system (1) is prohibitively expensive in terms of memory and inefficient to use.

Some ideas for preconditioning the minimisation of the cost function in the observations space, are: (i) the preconditioner proposed by Amodei (1995) that consists in scaling the linear system by the square root of the of observations error covariance matrix, see also Courtier (1997); (ii) the preconditioner of Egbert and Bennett (1996), see also Egbert (1997), that consists of building a preconditioning matrix by computing a few representer functions associated with pre-selected observations; (iii) the so-called second level preconditioner where a low rank approximation of the inverse of the matrix of the linear system is built with the iterations of the CG as in Giraud and Gratton (2006) and Chua et al. (2009). Two preconditioners are considered in this study: the preconditioner proposed by Amodei (1995), and the preconditioner proposed by Egbert and Bennett (1996). The convergence of these two preconditioners is compared using the representer based 4D-Var for the Navy Coastal Ocean model of Ngodock and Carrier (2014) for the assimilation of real observations of the ocean in the region around Hawaii. These two preconditioners have been successfully implemented in individual assimilation problems. However, our interest lies in their relative performance against each other when implemented within the same assimilation system. To our knowledge such comparisons have not been carried with realistic ocean models, although we note the work of Zaron (2006) who made similar comparisons using a planetary geostrophic model, and Chua et al. (2009) using an atmospheric model.

In the representer method of Bennett (1992, 2002), also known as the dual space formulation, Courtier (1997), the analysis equation for the optimal increment is written in the form

$$\delta x = LCL^{\mathrm{T}}H^{\mathrm{T}}(HLCL^{\mathrm{T}}H^{\mathrm{T}} + R)^{-1}(y - Hx^{\mathrm{b}}). \quad (1)$$

where C is the background error covariance for the entire model trajectory, L is the tangent linear model (TLM) and L^{T} is its adjoint, H is the linearised observation operator that projects the entire model trajectory onto the observations space. The solution of (1) is obtained by first solving the linear system

$$(HLCL^{\mathrm{T}}H^{\mathrm{T}}+\mathrm{R})\beta = (y - Hx^{\mathrm{b}}), \quad (2)$$

and applying the operator $LCL^{\mathrm{T}}H^{\mathrm{T}}$ (which is a succession of the adjoint integration, the covariance multiplication and the TLM integration) to the result. The most computationally intensive and expensive part of the assimilation consists of the iterative process for the inversion in (2), where two matrix-vector multiplications in observation space are needed at every iteration: the first, $R\beta$ is readily available since R is usually considered diagonal; the second, $HLCL^{\mathrm{T}}H^{\mathrm{T}}\beta$ is obtained by solving the adjoint model with impulse at observation locations scaled by the components of the vector β, applying the background and/or model error covariance and passing that information to the TLM, which in turn is evaluated at the observation locations. Although both matrices involved in the inversion in (2) are symmetric (the representer matrix $HLCL^{\mathrm{T}}H^{\mathrm{T}}$and the observation error covariance matrix R), making the use of efficient iterative algorithms such as the CG of Hestenes and Stiefel (1952) attractive, the system itself is usually poorly conditioned, thus worsening the computational expense of the 4D-Var algorithm.

The main problem that arises while solving the data assimilation problem in the observation space is that the commonly used CG algorithm does not generate iterates that monotonically reduce the quadratic cost function in the state space, El Akkraoui et al. (2008), Gratton and Tshimanga (2009) and El Akkraoui and Gauthier (2010). The First few iterations can lead to a solution that is worse than the background. To solve the problem of non-monotonic reduction of the quadratic cost function associated with the algorithm defined in the observation space, Gratton and Tshimanga (2009) derived an algorithm known as restricted preconditioned conjugate gradient (RPCG). The RPCG solves the data assimilation problem in the observation space and generates iterates that are mathematically equivalent to the iterates of the CG applied to the problem in the state space with the $\sqrt{\mathrm{B}}$ preconditioning. The RPCG relies on the existence of a special matrix (eqn. 8 in Gratton and Tshimanga (2009)) that is a function of the preconditioner used in the state space.

When the matrix of background error covariance is used as preconditioner for the primal problem, the special matrix is equal to the identity. This special case is known as the Restricted B-preconditioned Conjugate Gradient (RBCG) in Gürol et al. (2014), who also derived the Lanczos equivalent of the RBCG and called it

restricted B-preconditioned Lanczos (RBLanczos), and established the equivalence between the RBCG and the CG equipped with a suitable inner product to solve the dual problem preconditioned by R^{-1}. Examples were presented in Gürol et al. (2013) showing the benefits of RBCG in a 3D-VAR system for the global configuration of the Nucleus for European Modeling of the Ocean (NEMO) model, Madec (2008), and the benefits of RBLanczos in a 4D-Var system for a California Current configuration of the Regional Ocean Modeling System (ROMS) of Moore et al. (2011).

The preconditioners are briefly presented in Section 2, followed by experimental results in Section 3. The relative contributions of the observations error preconditioner and the representer matrix-based inner product are discussed in Section 4, and concluding remarks are given in Section 5.

2. The preconditioners

2.1. Amodei

Amodei (1995) proposed a preconditioner based on the observation error covariance matrix that transforms (2) into:

$$(R^{-1/2}HLCL^{\mathrm{T}}H^{\mathrm{T}}R^{-1/2}+\mathrm{I})\lambda = R^{-1/2}(y - Hx^{\mathrm{b}}), \quad (3)$$

and $\beta = R^{-1/2}\lambda$. The same CG algorithm used for solving (2) can be used for (3). Although (3) is technically just a change of variables from (2), it acts like a preconditioner in practice, allowing the iterative inversion to converge significantly faster (i.e. with fewer iterations) as shown below. In (3) the matrix multiplication $R^{-1/2}HLCL^{\mathrm{T}}H^{\mathrm{T}}R^{-1/2}\lambda$ follows the same steps described above, except for the additional scaling of the observation impulses by $R^{-1/2}$ both before the adjoint and after the TLM integrations.

There are two implementations of (3) in this study. The first, hereafter referred to as A1, is the regular CG algorithm to solve (3) using the Euclidian inner product. However, according to Gratton and Tshimanga (2009) the convergence of the CG can be further improved by replacing the Euclidian inner product with the one defined by the matrix $R^{-1/2}HLCL^{\mathrm{T}}H^{\mathrm{T}}R^{-1/2}$, the latter being symmetric and positive definite, Bennett (2002). This second implementation of (3) is referred to as A2.

2.2. Egbert–Bennett

Egbert and Bennett (1996) and Egbert (1997) proposed a preconditioner $\hat{P}$ (hereafter referred to as the EB preconditioner) as an approximation of the so-called stabilised representer matrix

$$P = HLCL^{T}H^{T} + R, \quad (4)$$

where a limited number of representer functions is computed and used to approximate the full representer matrix, and the combination of singular value decomposition and the Sherman–Morrison and Woodbury-formula are used to obtain the multiplication of $\hat{P}^{-1}$ with a vector in observation space. Since representer functions are independent from one another their computation can be done in parallel, reducing the overhead cost of the EB preconditioner. However, the pre-selection of the observation locations for which representer functions are computed for the EB preconditioner is attractive mainly for fixed observing platforms, for example, repeat period satellites, stationary buoys, and coastal high frequency radars. For moving platforms such as gliders, ships, or drifters (e.g. drifting buoys or Lagrangian drifters), a new selection has to be made in each assimilation cycle if observation locations from these instruments are included in the preconditioner. Should one choose not to include observation locations from moving platforms in the preconditioner, there would still be the problem of optimal distribution of the ones to include, as well as the number. Finally, the fact that representer functions are computed with linearised dynamics implies that the EB preconditioner needs to be recomputed in subsequent assimilation cycles because of the nonlinear variations of the background solution around which the dynamics are linearised. These portability issues are in stark contrast with the Amodei (1995) preconditioner which does not depend on the observation locations or the background solution, and thus is easily used from one assimilation cycle to next, regardless of the observation platforms.

3. The experiments

The model domain covers longitudes 162–153W and latitudes 17–24N, with a horizontal resolution of 6 km on a curvilinear grid of 157×130 points, while the vertical grid is composed of 35 sigma layers and 15 z-levels extending from the surface to the maximum depth of 5000 m. The model is forced with atmospheric fluxes fields (e.g. wind stress, IR radiation flux, etc.) from the Navy Operational Global Atmospheric Prediction System (NOGAPS) of Goerss and Phoebus (1992), and Rosmond et al. (2002) with 0.5° resolution, and initial and boundary conditions are taken from the global NCOM run at 0.125° resolution. The model is initialised on 1 June 2008 and allowed to spin-up until 1 July 2008.

An assimilation experiment is carried out in the region around Hawaii. Observations are processed in bins of 6 h duration, and consist of sea level anomalies (SLA) from satellite altimetry, sea surface temperature from satellite, and temperature and salinity profiles from expendable bathythermographs, Argo floats and a couple of gliders. One assimilation window of 10 days is chosen to allow a full cycle of sea surface height (SSH) from altimetry, in order to consider the SSH observation platform as fixed. There is a total of 753 SSH, 20,363 surface and subsurface temperature and salinity (SST), and 2818 combined subsurface temperature and salinity observations.

Observations coverage for the 10-day assimilation window is shown in Figure 1 for SSH, SST(SSS, profile T and profile S). Also shown in Figure 1 is the location of 100 observations selected for the EB preconditioner.

3.1. Residuals norm

The convergence of the CG algorithm is usually evaluated by the norm of the residual term relative to its initial value, that is, at iteration 0. This ratio is often used to stop the algorithm when it reached a pre-defined stopping criterion. Figure 2 shows the evolution of the norm of the residual in the CG algorithm. The relative residuals norm initially increases in the first few iterations with the EB preconditioning before it begins to slowly decrease, becoming smaller than its initial value after the 20th iteration. With A1, the relative residuals norm also increases initially for three iterations (in accordance with El-Akkraoui et al. (2008)), before it begins to decrease monotonically at a much faster rate than with the EB preconditioner. The A2 implementation has relative residuals norm that monotonically decrease from the first iteration and at a rate that is even faster than the A1. As an example, it can be seen from Figure 2 that it takes A2 14 iterations to reach the value 0.1; A1 reaches the same value in 17 iterations, and EB has not reached it in 90 iterations (not shown).

3.2. Sequence of analyses

After (2) or (3) is solved for the representer coefficients β, the optimal analysis increments are computed as

$$\delta x = LCL^{\mathrm{T}}H^{\mathrm{T}}\beta, \tag{5}$$

which is obtained by solving the adjoint model forced by Dirac impulses centred at the observations and scaled by the representer coefficient, convolving the adjoint solution with the background error covariance and using the convolved adjoint solution to force the TLM. The increments in (5) are then added to the background x^{b} to form the analysis. Thus, the analysis can be computed at each iteration of (3) to monitor the convergence of the algorithm not only at the observation locations, but also all over the domain, although the accuracy of the analysis can only be assessed at the observation locations. We will therefore evaluate the convergence of the CG algorithm with the preconditioners above in terms of the accuracy of the analysis per iteration, and the convergence of the analyses sequence.

3.2.1. Accuracy at the observation locations

The analysis error is defined here by a fit to the observations metric for each iteration as

$$J^k_{\mathrm{FIT}} = \frac{1}{M}\sum_{m=1}^{M}\frac{|y_m - H_m \mathrm{x}^k|}{\sigma_m}, \tag{6}$$

where y_m is the mth observation, M is the total number of observations, H_m is the observation operator, x^k is the assimilated solution or analysis after the kth iteration, and σ_m is the observation error or standard deviation.

The sequence defined by J_{FIT} in (6) is a good tool to monitor both the accuracy of the analysis at the observation locations and the convergence of the CG algorithm, because the analysis is expected to improve with the iterations. Figure 3 shows the evolution of the J_{FIT} sequence. With the EB preconditioning, the sequence shows a slightly chaotic pattern with a general decreasing tendency, and no apparent convergence after 30 iterations. With A1 there is an increase in the first iteration and a monotonic decrease from the second and subsequent iterations, becoming almost constant after 20 iterations. With A2 the sequence decreases monotonically from the first iteration and becomes almost constant after 12 iterations.

3.2.2. Accuracy everywhere

We consider the sequence of analyses (one per iteration) and examine the convergence through the sequence of differences between consecutive analyses normalised by the difference between the first analysis and the background solution.

$$e^k = \frac{\|x^k - x^{k-1}\|}{\|x^1 - x^b\|}, \quad k \geq 2, \tag{7}$$

where x^k is the entire analysis trajectory after the kth CG iteration, and x^{b} is the background. The sequence e^k should reveal the convergence of the analyses from one iteration to the next, even in regions that are distant from the observation locations.

From the evolution of the sequence e^k in Figure 4 it can be seen that there is no discernible pattern of convergence from the EB preconditioner after 30 iterations, that

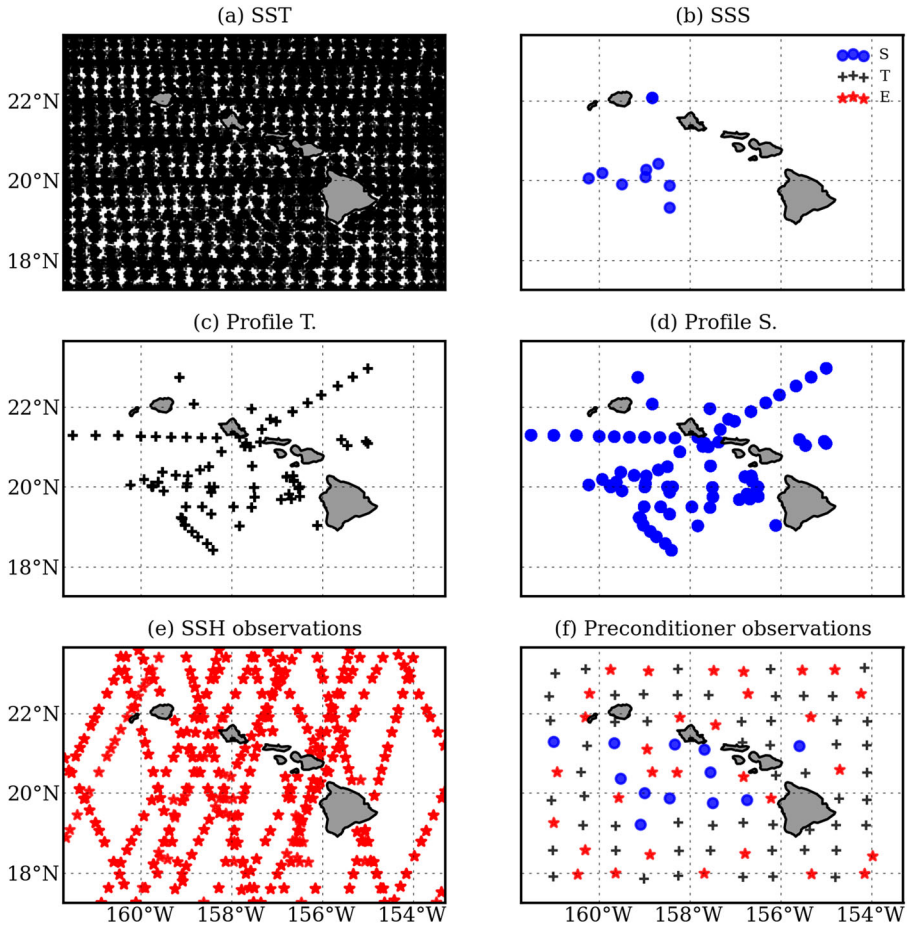

Figure 1. Observation distribution in the model domain during the 10-day assimilation window: SST (a), SST (b), temperature profiles (c), salinity profiles (d), SSH (e), and the observations selected for the construction of the EB preconditioner (f).

is, the consecutive analyses are quite different from one another. With A1, the amplitude of the correction shows an increasing tendency in the first six iterations followed by a general decreasing tendency. The increasing tendency in the first six iterations suggests that the correction is more important from one iteration to the next. This is not the case with A2, where the sequence decreases monotonically from the first iteration. Compared to Figures 2 and 3, Figure 4 shows that while convergence may be reached in 12–14 iterations for A2 (17–20 iterations for A1) in the observations space as measured by either the residual norm or the fit to the observations, the sequence of analysis differences in the state space converges at a slower rate.

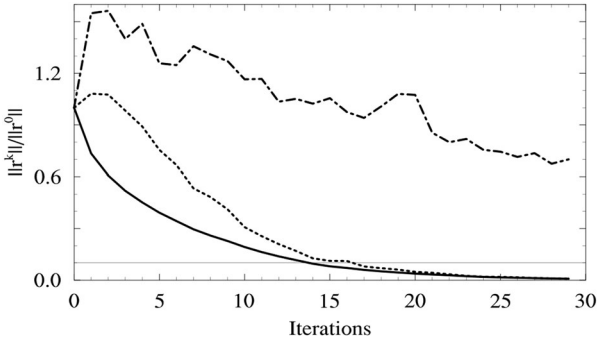

Figure 2. Evolution of the relative norm of the residuals with the CG iterations using the EB preconditioner (dash-dotted line), and the A1 (dotted line) and A2 (solid line) implementations of the Amodei preconditioner.

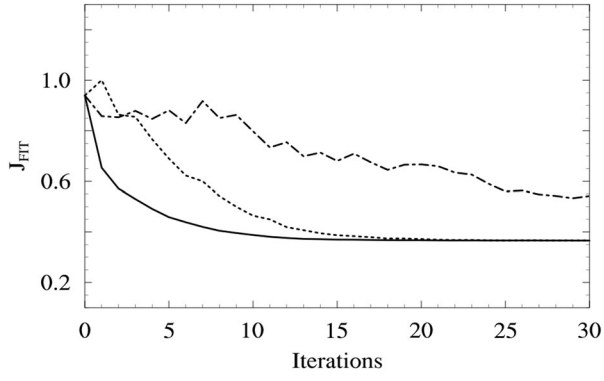

Figure 3. The convergence of the analysis error with CG iterations from the BE (dash-dotted line) and the A1 (dotted line) and A2 (solid line).

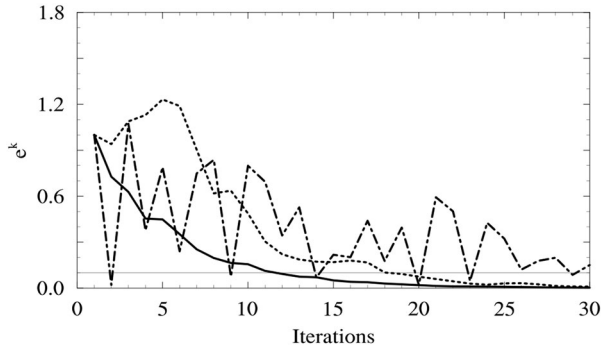

Figure 4. Evolution of the sequence e^k with the iterations of the CG algorithm using the EB (dash-dotted line), the A1 (dotted line) and the A2 (solid line) preconditioners.

3.3. Gradient of the cost function

The convergence of the analysis (i.e. the entire model trajectory) cannot be tested at each iteration because of the computational burden of extracting the analysis, which requires the integrations of the adjoint and TLM. However, the behaviour of the analysis convergence is assessed through the behaviour of the gradient of the cost function in the state space. The norm of the gradient of the cost function in the state space can be computed at no additional cost thanks to the re-organisation of the computations in the CG algorithms as shown by Gratton and Tshimanga (2009). Figure 5 shows the evolution of the gradient of the cost function in the state space. The behaviour matches that of the evolution of the solution in the state space, Figure 3. In the case of the Euclidean inner product, the gradient is larger than the initial gradient for the first six iterations, it then decreases monotonically. In the case of the modified inner product, there is a monotonic decrease from the first iteration.

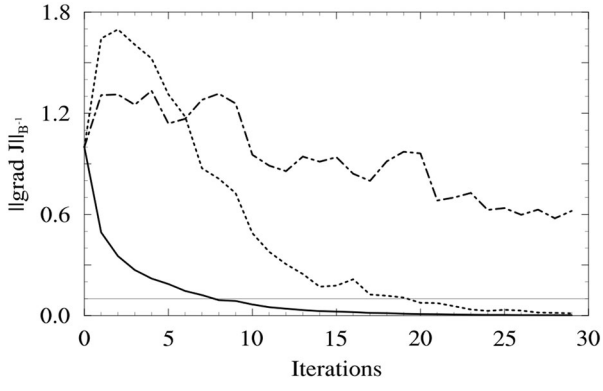

Figure 5. Evolution of the norm of the gradient of the cost function with iterations of the CG algorithm using the EB (dash-dotted line), the A1 (dotted line) and the A2 (solid line) preconditioners.

4. Preconditioner or inner product?

Results from the experiments above clearly indicate that the best performance of the CG in the minimisation of the cost function is achieved with the combination of the Amodei preconditioner and the inner product defined by the representer matrix. In an attempt to elucidate the relative contribution of each of these two factors, that is, the Amodei preconditioner and the inner product, two additional experiments are carried out with the non-preconditioned CG: one uses the Euclidian inner product (hereafter referred to as N1) and the other uses the inner product defined by the representer matrix (hereafter referred to as N2). The results of these experiments are compared to those from A1 and A2 experiments, using all four convergence criteria above.

Figure 6 shows that common conclusions can be drawn from all four metrics. For the often used criteria such as the residual and gradient norms there is no apparent convergence pattern for N1. Although its J_{FIT} values decrease and seem to be reaching an asymptote, they still show some variability from one iteration to another as seen with the sequence of analysis differences e^k.

However, the same cannot be said of N2. The latter shows a decrease in the gradient norm that is similar to A2 for all iterations. For the residual norm the decrease in N2 is similar to A2 in the first nine iterations, after which N2 continues to decrease at a slower rate than A2, but still faster than A1 until the 16th iteration. Also, N2 shows a better decrease in the J_{FIT} and e^k metrics to the 9th and 17th iterations respectively. Finally, except for the gradient norm, all other three metrics show that A1 has lower values than N2 in later iterations. It can thus be said that the inner product defined by the representer matrix plays a significant role in the early iterations of the CG, but that inner product alone is not sufficient to achieve the desired convergence. The Amodei preconditioner helps the CG algorithm converge even further in the later iterations. Overall, the inner product defined by the representer matrix has a strong effect for the first range of iterations (10–15 for this experiment) and the preconditioning effect is predominant after that. However, only their combined effect ensures the monotonic decrease of all the studied convergence criteria.

It should be noted that the implementation of the modified inner product with the Amodei preconditioner involves only one additional application of the matrix-vector product, to initialise to iterative process. This is a rather marginal cost, equivalent to one additional iteration in the process, that significantly impacts the

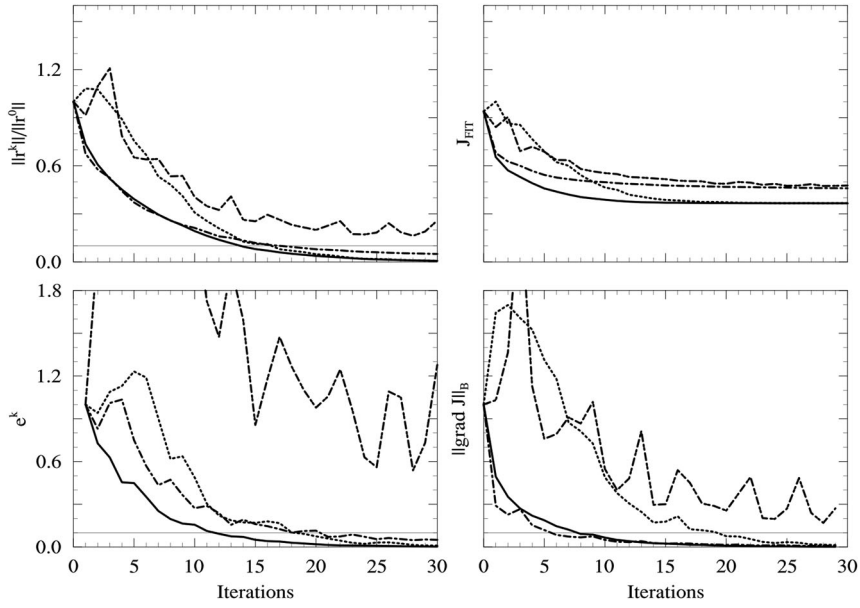

Figure 6. A comparison of A1 (dotted line), A2 (solid line), N1 (dashed line) and N2 (dash-dotted line) using the residual norm (top left), the fit to the observations (top right), the sequence of analysis differences (bottom left) and the norm of the gradient (bottom right) metrics.

minimisation as seen above: it was shown to stabilise and speed the convergence by many iterations depending on the convergence criterion.

5. Discussion and conclusions

Minimising the 4D-Var cost function in the observations space requires the solution of a linear system that is poorly conditioned, usually with iterative algorithms. When symmetric matrices are involved, as is the case with the representer method, CG methods prove to be very efficient. The assimilation of ocean observations using the representer based 4D-Var for NCOM was considered, along with two methods of preconditioning the linear system: the scaling of the linear system by the square root of the observation error variances from Amodei (1995), and the approximated inverse of the stabilised representer matrix based on the computation of some representer functions by Egbert and Bennett (1996). The former had two implementations of the CG, one with the Euclidian inner product, and the other with the inner product defined by the scaled representer matrix.

The three implementations of the CG were compared in a 10-day assimilation experiment in the region around Hawaii, involving almost 25,000 observations. The comparison was based on four criteria: the norm of the residuals in the linear system, the analysis error evaluated at the observation locations, the convergence of the sequence of analyses, and the norm of the gradient of the cost function. All four criteria show that in the first 30 iterations of the CG the A2 implementation provided the fastest convergence, followed by A1. There was a general monotonic decrease in every criterion for A2; in A1 the criteria tended to increase in the first few iterations before they began a general decrease, whereas EB did not show a discernible pattern of convergence. Convergence was reached with A2 (A1) in as few as 14(17–20) iterations for the residuals norm and the analysis errors, and 21(28) and 25(29) iterations for the gradient norm and the sequence of analyses respectively. Note that each iteration of the CG has a computational cost that is equivalent to that of one representer function. Thus the convergence numbers above for both A2 and A1 show significant savings in computational costs compared to the computation of the 100 representer functions used in the EB implementation, especially given that the latter did not converge after 30 iterations.

The relatively inferior performance of the EB preconditioner may be attributed to a few factors. The first is linearity: we are dealing with a 50-layer nonlinear hydrostatic model while the original EB preconditioner was applied to a quasi-linear two-dimensional tide model. Second, a rather small number of representer functions (100 out of 25,200 observations), compared to a significantly large number of representer functions were computed (10^4 out of 10^5 observations in Egbert et al. (1994) or 1200 out of 6000 observations in Egbert (1997)). Finally, the observing network had a quasi-uniform coverage of the global tide model domain in Egbert et al. (1994), a coverage that is nearly impossible to achieve with the observations at hand, especially in the ocean subsurface.

In the absence of an appropriate inner product, the first few iterations can deteriorate the initial approximation and the algorithm can require more iterations, even with the implementation of a preconditioner. It was found that the inner product defined by the representer matrix played a significant role in the earlier iterations, evidenced by a rapid decrease in all four convergence metrics. However, with that inner product alone the CG algorithm converged towards a solution that was less accurate than the solution obtained from the CG equipped with the Amodei preconditioner, according to the fit to the observations metric. Thus, for the minimisation of the data assimilation cost function in the observations space, the implementation of the appropriate inner product in the CG algorithm should be done in conjunction with the Amodei preconditioner. Neglecting the latter would yield a misleading sense of convergence, especially if a small number of iterations are required.

Acknowledgment

The authors would like to thank the anonymous reviewers for their constructive comments that helped to improve the quality of the manuscript.

Disclosure statement

No potential conflict of interest was reported by the authors.

Funding

This work was sponsored by the Office of Naval Research Program Element 62435N as part of the "NCOM-4DVAR" and the "A multiscale Approach for Assessing Predictability of ASW environment" projects. This paper is NRL paper contribution number NRL/JA/7320-15-2472.

References

Amodei L. 1995. Solution approchée pour un problème d'assimilation de données météorologiques avec prise en compte de l'erreur de modèle. Comptes Rendus de l'Académie des Sciences. 321:série II a, 1087–1094.

Bennett AF. 1992. Inverse methods in physical oceanography. New York: Cambridge University Press. 347 pp.

Bennett AF. 2002. Inverse modeling of the ocean and atmosphere. Cambridge: Cambridge University Press.

Chua BE, Xu L, Rosmond T, Zaron ED. 2009. Preconditioning representer-based variational data assimilation systems: application to NAVDAS-AR. Data Assimilation Atmos, Oceanic Hydrol Appl. 2009:307–319.

Courtier P. 1997. Dual formulation of four-dimensional assimilation. Q J R Met Soc. 23:2449–2461.

Egbert GD. 1997. Tidal inversion: interpolation and inference. Progr Oceanogr. 40:53–80.

Egbert GD, Bennett AF. 1996. Data assimilation methods for ocean tides. In Malanotte-Rizzoli P. editor. Modern approaches to data assimilation in ocean modeling. New York, NY: Elsevier; p. 147–179.

Egbert GD, Bennett AF, Foreman MGG. 1994. TOPEX/POSEIDON tides estimated using a global inverse model. J Geophys Res. 99:24821–24852. doi:10.1029/94JC01894

El Akkraoui A, Gauthier P, Pellerin S, Buis S. 2008. Intercomparison of the primal and dual formulations of variational data assimilation. Q J R Meteorol Soc. 134:1015–1025. doi:10.1002/qj.257

El Akkraoui A, Gauthier P. 2010. Convergence properties of the primal and dual forms of variational data assimilation. Q J R Meteorol Soc. 136:107–115.

Giraud L, Gratton S. 2006. On the sensitivity of some spectral preconditioners. SIAM journal on matrix analysis and applications 27:1089–1105.

Goerss JS, Phoebus PA. 1992. The Navy's operational atmospheric analysis. Wea Forecast. 7:232–249.

Golub GH, Van Loan CF. 2013. Matrix computation. 4th ed. Baltimore (CA): The Johns Hopkins University Press. pp 65.

Gratton S, Tshimanga J. 2009. An observation-space formulation of variational assimilation using a restricted preconditioned conjugate gradient algorithm. Q J R Meteorol Soc. 135:1573–1585. doi:10.1002/qj.477

Gürol S, Weaver AT, Moore AM, Piacentini A, Arango HG, Gratton S. 2014. B-preconditioned minimization algorithms for variational data assimilation with the dual formulation. Q J R Meteorol Soc. 140:539–556. doi:10.1002/qj.2150

Hestenes MR, Stiefel, E. 1952. Methods of conjugate gradients for solving linear systems. J Res Nat Bureau Stand. 49:409–436.

Lorenc AC. 1988. Optimal nonlinear objective anlaysis. Q J R Meteorol Soc. 114:205–240. doi:10.1002/qj.49711447911

Madec G. 2008. NEMO ocean engine. Technical Report 27, Note du P^ole de modélisation, Institut Pierre-Simon Laplace (IPSL), France. Available from: http://www.nemo-ocean.eu.

Moore AM, Arango HG, Broquet G, Powell BS, Zavala-Garay J, Weaver AT. 2011. The Regional Ocean Modeling System (ROMS) 4-dimensional variational data assimilation systems. Part I: system overview and formulation. Prog Oceanogr. 91:34–49.

Ngodock H, Carrier M. 2014. A 4dvar system for the Navy Coastal Ocean model. Part I: system description and assimilation of synthetic observations in Monterey Bay. Mon Wea Rev. 142:2085–2107. doi:http://dx.doi.org/10.1175/MWR-D-13-00221.1

Robert C, Blayo E, Verron J. 2006. Reduced-order 4D-Var: a preconditioner for the full 4dvar data assimilation method. Geophys Res Lett. 33:L18609. doi:10.1029/2006GL026555

Rosmond TE, Teixeria J, Peng M, Hogan TF, Pauley R. 2002. Navy operational global prediction system (NOGAPS): forcing for ocean models. Oceanography. 15:99–106. doi:10.5670/oceanog.2002.40

Saad Y. 2003. Iterative methods for sparse linear systems. 2nd ed. Society for Industrial and Applied Mathematics, ISBN-13: 978-089871534.

Zaron ED. 2006. A comparison of data assimilation methods using a Planetary Geostrophic model. Mon Wea Rev. 134:1316–1328. doi:10.1175/MWR3124.1