

Forecasts from fits of frontal fluctuations ¹

Arthur J. Miller ^{*}, Bruce D. Cornuelle

Climate Research Division, Scripps Institution of Oceanography, La Jolla, San Diego, CA 92093-0224, USA

Received in revised form 4 December 1998; accepted 14 January 1999

Abstract

A primitive equation ocean model is fit with strong constraints to non-synoptic hydrographic surveys in an unstable frontal current region, the Iceland–Faeroe Front. The model is first initialized from a time-independent objective analysis of non-synoptic data (spanning 2 to 6 days). A truncated set of eddy-scale basis functions is used to represent the initial error in temperature, salinity, and velocity. A series of model integrations, each perturbed with one basis function for one dependent variable in one layer, is used to determine the sensitivity to the objective-analysis initial state of the match to the non-synoptic hydrographic data. A new initial condition is then determined from a generalized inverse of the sensitivity matrix and the process is repeated to account for non-linearity. The method is first tested in ‘identical twin’ experiments to demonstrate the adequacy of the basis functions in representing initial condition error and the convergence of the method to the true solution. The approach is then applied to observations gathered in August 1993 in the Iceland–Faeroe Front. Model fits are successful in improving the match to the true data, leading to dynamically consistent evolution scenarios. However, the forecast skill (here defined as the variance of the model–data differences) of the model runs from the optimized initial condition is not superior to less sophisticated methods of initialization, probably due to inadequate initialization data. The limited verification data in the presence of strong frontal slopes may not be sufficient to establish forecast skill, so that it must be judged subjectively or evaluated by other quantitative measures. © 1999 Elsevier Science B.V. All rights reserved.

Keywords: Ocean forecasting; Iceland–Faeroe Front; Data assimilation; Ocean modeling

^{*} Corresponding author. E-mail: ajmiller@ucsd.edu

¹ These results were presented at the 1998 AGU Ocean Sciences Meeting in the Tribute to Allan Robinson Session under the title ‘*In Vino Veritas et in Veritas Vino: Optimizing HOPS Skill, Italian Style*’.

1. Introduction

Ocean mesoscale eddy forecasting is typically limited by a number of factors, including inadequate initialization information, unknown boundary conditions, inaccurate model physics, and atmospheric forcing functions that must also be predicted. Moreover, establishing skill levels in mesoscale forecasting is also limited by inadequate validation data and the ambiguity of defining skill. It is of interest to explore the consequences of these issues to aid the development of mesoscale ocean forecasting techniques such as have been developed by the Harvard team (e.g., Robinson, 1996).

The region around the Iceland–Faeroe Front (IFF, hereinafter) is vigorously unstable with rapidly evolving small-scale eddies and frontal meanders that have time scales as short as 2 days and length scales as small 10 km (e.g., Hopkins, 1991; Niiler et al., 1992; Perkins, 1992; Allen et al., 1994; Tokmakian, 1994; Miller et al., 1995a,b; Poulain et al., 1996). A unique dataset for mesoscale forecasting experiments was collected in the IFF during August 1993. The data contains a relatively finely resolved initial hydrographic survey, an updating survey, and a final verification survey (Fig. 1). This data was used to make real-time forecasts at sea and the results have been evaluated quantitatively for forecast skill (pattern correlation and rms error) for both a primitive equation model (Robinson et al., 1996a) and a quasigeostrophic model (Miller et al., 1995b).

Although the skill scores for those sets of forecasts were encouraging, there are many issues that can be explored with a dataset designed specifically for forecasting and validation. For example, the Initialization Survey was collected over a 3-day time interval which poses a problem with choosing synoptic initial conditions. Robinson et al. (1996a) used a feature model strategy combined with optimal interpolation to launch forecasts from the end of the 3-day survey. Can a more objective initialization strategy that allows for this non-synopticity improve the skill of forecasts of independent data?

In this paper, we examine the issues of initialization and verification of IFF ocean forecasts by applying an inverse method for initialization in two sets of experiments. The inverse method is similar to Bennett's representor method (Bennett, 1992) and Wunsch's Green's function approach (Wunsch, 1996) and enforces the dynamics as a strong constraint. We adjust only the most energetic scales of the model initial conditions to minimize the model–data misfit variance in the fitting time interval. This truncation is related to the ensemble Kalman filter discussed by van Leeuwen and Evensen (1996) and by Lermusiaux (1997) and Lermusiaux et al. (1998) for the same Harvard primitive equation model used here. We first test the inverse method by generating synthetic data, sampled in the same way the observations were collected, in 'identical twin' predictability experiments. We then test our techniques on the observations to determine if the fits are successful and if forecast skill is enhanced over the results of Robinson et al. (1996a).

We find that we are able to fit the model to the non-synoptic hydrographic data even for week-long time intervals. This result suggests that the model physics is acceptably consistent with the flow evolution during the August 1993 IFF surveys. However, increasing the skill of the fits does not always yield increased forecast skill, at least when gauging skill by rms error variance. This is due to the initialization data being too

AUGUST 1993 CRUISE HYDROCAST DATA

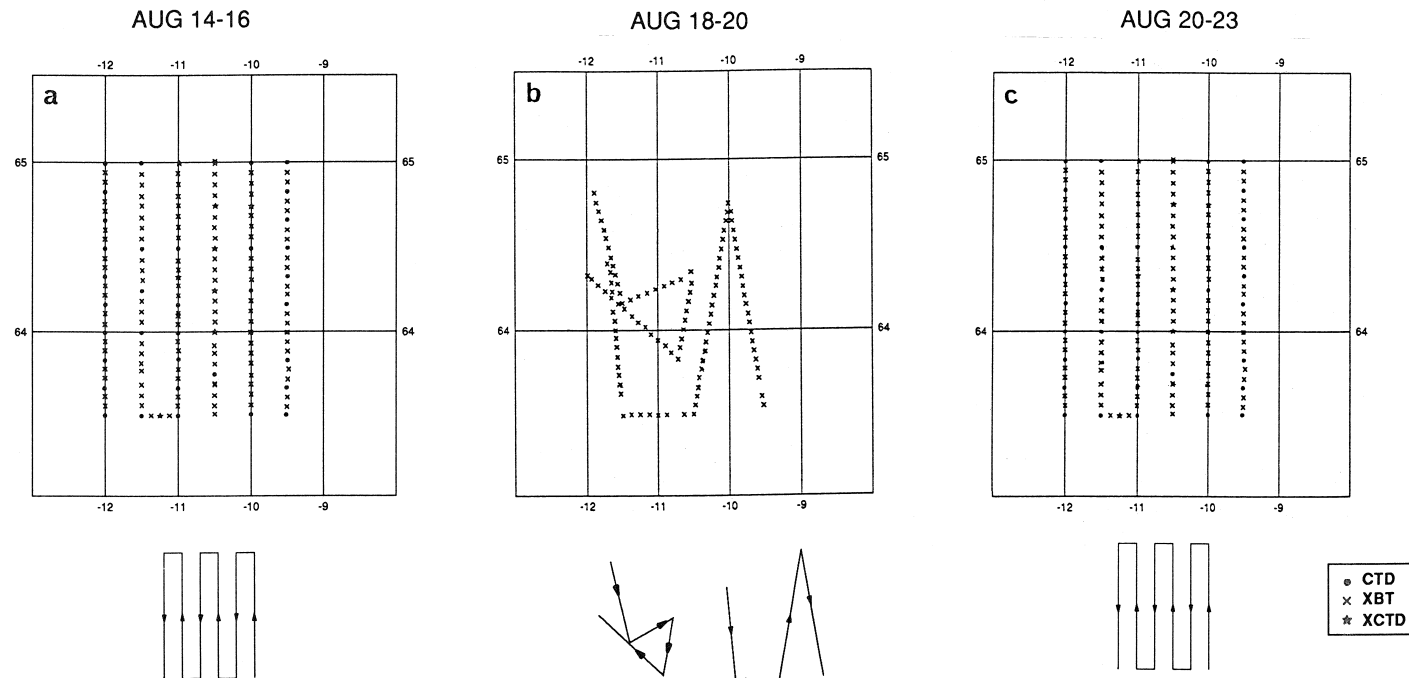


Fig. 1. Maps of the locations of the three hydrographic surveys collected in the Iceland–Faeroe Front region during August 1993. (a) Initialization Survey (approximately 3 days duration), (b) Zig-Zag Survey (approximately 2 days duration) and (c) Validation Survey (approximately 3 days duration). The type of data is indicated by the symbol: dot for CTD, cross for XBT and star for XCTD. The path followed by the ship is indicated below each plot.

limited to bring the model close enough to the true initial condition in a region which is very sensitive to initial condition error. Moreover, the limited verifying data does not always give an unambiguous demonstration of quantitative forecast skill. Qualitative measures of assessing fitting and forecasting skill must then be invoked to judge the results.

We first introduce the data (Section 2), the model (Section 3), and the fitting technique (Section 4), then test the method in an identical twin experiment (Section 5). We then apply the approach to the August 1993 IFF dataset (Section 6) and then summarize and discuss the results (Section 7).

2. Initialization and validation data

In August 1993, hydrographic surveys were collected in the region around the IFF in a joint project between the SACLANT Undersea Research Center and Harvard University aboard the NATO vessel R/V Alliance. These surveys were specifically designed for mesoscale eddy forecasting experiments in that they included an Initialization Survey, an (updating) Zig-Zag Survey, and a Validation Survey (Fig. 1), hereinafter denoted with capital letters for clarity. The hydrographic data consist of expendable bathythermograph (XBT), conductivity–temperature–depth (CTD), and expendable CTD (XCTD) data as indicated by the symbols in Fig. 1. The Initialization Survey data spanned August 14–16 and included XBTS, CTDs, and XCTDs sampled at 24 km zonal resolution and 7 km meridional resolution. The Zig-Zag Survey, from August 18 to 19, included only XBT data in and around the most actively evolving area of the IFF at that time. The Validation Survey followed the same track as the Initialization Survey from August 20–24. Additional data was also collected, including surface drifter displacements in the core of the IFF, one clear satellite image of sea-surface temperature on August 22, and subsurface velocity measurements from moored currents along 12° W. Miller et al. (1995b), Miller et al. (1996), and Robinson et al. (1996a) provide additional details of these data and their acquisition.

To make the first guess for the model initialization and to provide a concise field picture of the data, the hydrographic data were objectively mapped as if they were time-independent. All measurements from a particular survey were used together to produce smooth, three-dimensional temperature and salinity fields (Fig. 2).

A smooth background state for the objective analysis (OA, hereinafter) was constructed from the observations and historical data, using the historical data to determine the background state for the portion of the model domain outside the observations. The background state was mapped independently at each depth level by fitting the temperature and salinity fields to two-dimensional quadratic polynomials, which approximated the frontal region as a transition between relatively homogeneous water masses to the north and south.

The anomalies with respect to the background state were mapped objectively (Bretherton et al., 1976), using a Gaussian horizontal covariance with an e-folding scale of 20 km. The vertical smoothing was done by least-squares fitting each temperature station to three vertical empirical orthogonal functions (EOFs, hereinafter) derived from

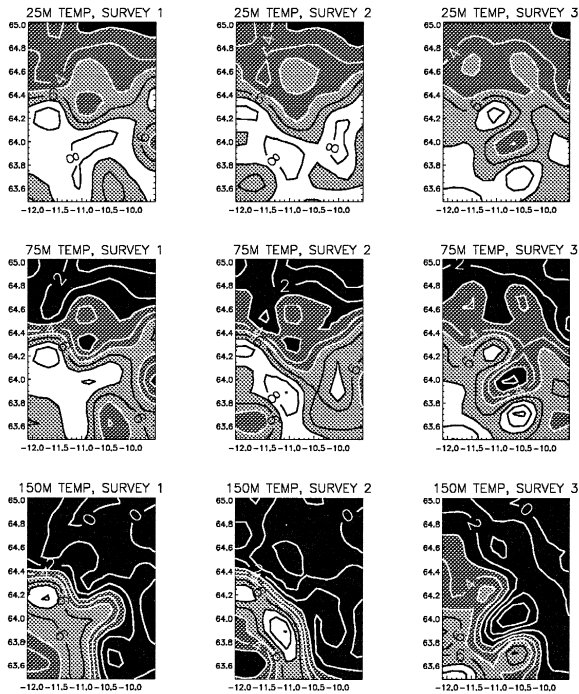


Fig. 2. Objective analyses of the vertical-EOF-filtered temperature at three depths (25, 75, 150 m) corresponding to the top three model surfaces for (left) Initialization Survey, (center) Zig-Zag Survey plus Initialization Survey, and (right) Validation Survey. Temperature warmer than 7°C is white, warmer than 5°C light grey, warmer than 3°C dark grey and cooler than 1°C black. Contour interval is 1°C . These fields are used for the first guess at the model initial conditions and for subjectively evaluating the flow fields. Model–data misfit is always computed using raw unfiltered data.

the entire temperature dataset, and then the EOF amplitudes were mapped horizontally. Salinity was supplied for XBT stations using T – S correlations derived for the EOF amplitudes. The three temperature EOFs explained 96% of the observed temperature variance, using uniform weighting in the vertical. Put another way, the stations reconstructed from three temperature EOFs differed from the original stations by about 0.5° rms at the surface, decreasing linearly to about 0.3° rms at 360 m. The T – S conversion was checked on the 53 CTD casts, and calculated salinities differed from observed salinities with a mean of less than 0.005, and rms of 0.07 psu above 60 m and 0.04 psu from 60–360 m. Below 360 m, there was very little salinity variation in the observations.

Fig. 2 shows the results of the temperature OAs for the three surveys plotted to show the general nature of the evolving frontal structures. Initially, the front has a kink in its east–west path, with an apparent eddy lying to the north of the front. Two days later, the front exhibits a southeastward flow in the western part of the domain. Roughly 4 days later, a hammerhead instability develops in the front, clearly seen in a satellite image of SST (Miller et al., 1995b; Robinson et al., 1996a). These general and qualitative features

of the frontal evolution will prove useful in verifying the fits and forecasts of the model, since quantitative skill scores may be misleading as discussed below. *N.b.*, only the raw data (with uncertainty bounds) are used to quantitatively validate the model; the OAs are used for initialization and qualitative discussions only.

3. Primitive equation forecast model

The Harvard Ocean Prediction System (HOPS) has been discussed in the context of many uses around the world (e.g., Robinson, 1996; Robinson et al., 1996b, 1998) Here we summarize only the rudiments of the primitive equation (PE, hereinafter) model version 7.28 used for this study. The Geophysical Fluid Dynamics Laboratory PE model (Bryan and Cox, 1967; Semtner, 1974) provides the basic integration algorithm for the rigid-lid, open-boundary PE model of HOPS. The PE open boundary conditions, subgrid scale physics, and bottom topographic treatment via hybrid coordinates were developed at Harvard (Spall and Robinson, 1989; Lozano et al., 1994).

In this study, we use a grid similar to that employed by Robinson et al. (1996a) in real-time forecasts. The horizontal resolution is 5 km in both directions. The hybrid coordinates are used in such a way that the top three levels are flat while the bottom two levels follow the strongly varying topography (which is always shallower than the flat levels). The five surfaces of vertical resolution consist of the top three level surfaces at 25, 75, and 150 m depth and the bottom two sigma coordinate surfaces spaced equally between 200 m and the topographic ridge in the area (Fig. 3). The topography was clipped to 300 m depth (including rendering Iceland as a shallow sea) then smoothed

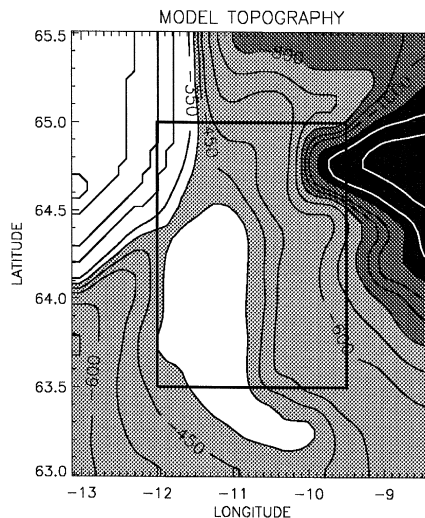


Fig. 3. Topography of the region around Iceland and the Faeroe Islands used by the model. The ETOPO5 topography (including Iceland itself) was first clipped to 300 m minimum depth, then smoothed several times with a grid scale filter. Depths shallower than 400 m are white, less than 800 m light grey, less than 1200 m dark grey, and deeper than 1200 m black. Contour intervals are irregularly spaced to guide the eye. The data domain is indicated by the black rectangle while the model boundaries encompass the entire plot.

several times at grid scale to arrive at the final ‘conditioned’ topography. The horizontal boundaries of the model domain extend far (compared to Robinson et al., 1996a) from the data area in order to attempt to avoid the problem of specifying persistent or predicted boundary conditions. However, there is a trade-off on including the necessary extra grid points for distant boundaries vs. making the perturbation runs more computer time-efficient. The far-field boundary conditions are no-flux on temperature and salinity, Spall and Robinson (1989) radiation conditions on velocity and vorticity, and Orlanski (1976) radiation conditions on barotropic streamfunction. Although open boundary conditions in PE models are generally ill-posed (Bennett, 1992), they apparently have little influence on the flow in the data region during the rather short 3–10 day integrations considered here. Horizontal friction is modeled with a fourth-order Shapiro filter applied eight times each time step. Vertical diffusion has a constant coefficient of $5 \text{ cm}^2/\text{s}$.

A fundamental assumption in these experiments is that intrinsic variability of the IFF current and eddy field is the dominant mechanism generating variability. The effects of surface wind, heat flux, and fresh-water flux forcing are ignored (no-flux surface boundary conditions). The model eddy field evolves in a manner that is dynamically consistent (i.e., follows the equations of motion with no unphysical forcing during the run) after the specification of the initial conditions and the boundary conditions. Except for some initial tuning of parameters to give reasonably energetic and smooth flow conditions, we did not explore the sensitivity of these results to changes in model parameters, topographic smoothness or surface forcing. Each of these effects could be important and should be explored in future studies.

Model runs are initialized from OA (Section 2) of temperature and salinity computed on constant depth levels and mapped to the model surfaces. Baroclinic velocities are computed geostrophically from the temperature and salinity profiles with a level of no motion at 350 m depth. We tested the efficacy of holding the temperature and salinity fixed while integrating the model for several days to allow the velocity fields to come near to a dynamical equilibrium (cf. Bogden et al., 1993). However, this resulted in energetic initial conditions that rapidly produced evolving temperature and salinity fields that differed greatly from the observations in the model fitting time interval and increased the non-linearity of the attempts to fit the data. We thus chose to use the level of no motion velocity initialization with the inverse solutions giving the corrections to this guess of the initial state.

4. Inverse initialization strategy

Many techniques exist for combining data with models (e.g., see the reviews by Ghil and Malanotte-Rizzoli, 1991; Bennett, 1992; and Robinson et al., 1998). Least-squares methods are widely used for fitting both steady and unsteady models to data, and can be implemented sequentially as the Kalman filter followed by a smoother, or globally by solving the Euler–Lagrange equations to find the minimum of an objective function (Le Dimet and Talagrand, 1986; Scheinbaum and Anderson, 1990; Bennett, 1992; Bogden et al., 1996). The objective function is a sum of quadratic terms penalizing misfit between

the observations and the data generated by the model, and also penalizing corrections to the assumed model parameters, which may include forcing, initial, and boundary conditions. The weighting of the penalty terms often includes smoothness criteria. The error in the model equations can be thought of as time-dependent, non-physical forcing at every grid point.

In our case, the model initial state was the only adjustable field, and was adjusted to minimize J , the weighted sum of squared data misfits plus the weighted sum of initialization adjustments

$$J = \Psi^T P_{\Psi}^{-1} \Psi + [\mathbf{O} - \mathbf{F}(\Psi + \Psi_0)]^T R^{-1} [\mathbf{O} - \mathbf{F}(\Psi + \Psi_0)] \quad (1)$$

where Ψ is the adjustment to the initial model state vector, Ψ_0 ; P_{Ψ}^{-1} is the inverse of the covariance of the uncertainty in Ψ_0 ; \mathbf{O} is the vector of observations; $\mathbf{F}()$ is the vector of model variables, sampled at times and locations corresponding to \mathbf{O} ; R^{-1} is the inverse of expected covariance for the residuals; and $()^T$ denotes matrix transpose. By ‘data misfits’ we mean the differences between the observations and the model estimate at every data point, $d = \mathbf{O} - \mathbf{F}(\Psi + \Psi_0)$. The weighting of these misfits in the objective function is set by R^{-1} , which is assumed to be diagonal. The residuals are the components of the misfits which cannot be fit by adjusting the model initialization. They may be due to measurement errors, unmodeled structure, or incorrect model physics. The square root of the diagonal of matrix R is the vector of expected rms residual at the data points. Since we refer to this frequently, we will use the shorthand expression σ , where the elements are $\sigma_k = \langle r_k^2 \rangle^{1/2}$ (the square roots of the diagonal elements of R), indexed by datum.

The adjustments, Ψ , to the initial model state, Ψ_0 , are weighted by the inverse of the covariance of uncertainty in the initial conditions, P_{Ψ} , which is not diagonal because large-scale errors are expected to be more energetic than small-scale errors, although errors are assumed to be uncorrelated between layers and between variables (T and S uncorrelated, for example). This covariance function was generated for each layer from fields of rms initial condition uncertainty and horizontal decorrelation scale at each level. The rms uncertainty was derived from the temperature gradient magnitude for each level in the objectively mapped initialization, smoothing with a Gaussian filter with a 30-km e-folding scale, in order to spread the variance over the general frontal region. The resulting assumed rms uncertainty varied by a factor of 10 between the frontal region and the far-field. This has the effect of concentrating corrections in the frontal region, where the initialization uncertainty is highest. The decorrelation scales were set inversely proportional to the uncertainty, so that points near the front had a decorrelation scale of about 20 km, increasing to 80 km for points far from the front. This allows for small-scale corrections near the front and larger-scale corrections away from it. The model initialization included temperature, salinity, baroclinic current, and barotropic streamfunction, and the same covariance function was assumed for all variables at each level, although the variance levels were scaled to match the observed variability. This is an ad hoc choice for generating the covariance function, but we did not find a strong sensitivity to choices of smoothing or decorrelation scale.

To reduce the size of the problem, the state vector uncertainty covariance matrix, P_{Ψ} , was factored into eigenvalues and eigenvectors, and only the eigenvectors corre-

sponding to the largest 100 eigenvalues for each layer were used as ‘basis functions’ for the perturbations to model initial conditions. These accounted for 79 to 89% of the assumed uncertainty variance. This approach is frequently used in ensemble Kalman filter methods (e.g., van Leeuwen and Evensen, 1996). The twin experiments provided a check on the efficacy of these functions for representing error in an initial state; initial state temperature errors were well-represented although velocity errors were less so. Customized basis functions for velocity (based on velocity gradients, etc.) were not used. The factorization of P_{Ψ} simplified the penalty term for the perturbations to the initial conditions to:

$$\Psi^T P_{\Psi}^{-1} \Psi \approx \mathbf{m}^T P^{-1} \mathbf{m} \quad (2)$$

where \mathbf{m} is the vector of amplitudes of the basis functions and P is the diagonal matrix of eigenvalues for the retained basis functions. Note that this reduces the dimensionality of the initial condition adjustment but does not change the state vector of the dynamical model.

To compute the dependence of data misfit on basis function amplitude (‘model sensitivity’), a series of perturbation experiments was executed. In each experiment, the model was perturbed by changing the initial conditions using a single basis function. The model was run from the new initialization throughout the time with data, saving the differences between the predicted data values from the original and new initializations. The dependence of data on initial conditions is thus linearized around the current best guess of the model initialization. Changes in output model predictions for the data misfits, \mathbf{d} , are related to changes, \mathbf{m} , in the unknowns and the residuals, \mathbf{r} . The dependence is written in matrix notation as:

$$\mathbf{d} = H\mathbf{m} + \mathbf{r}, \quad (3)$$

where H is the matrix of derivatives (the linearized ‘forward problem’) and $\langle \mathbf{r}\mathbf{r}^T \rangle = R$.

Using the linearization and the orthogonal basis functions, the objective function to be minimized is:

$$J = \mathbf{m}^T P^{-1} \mathbf{m} + (\mathbf{d} - H\mathbf{m})^T R^{-1} (\mathbf{d} - H\mathbf{m}). \quad (4)$$

This seeks to minimize the weighted misfits between the observations and the model as well as the weighted size of the correction to the initial state. The weighting matrices P and R are now both diagonal, greatly simplifying the computations.

The model parameters which minimize the objective function are determined by iterating a linearized, time-dependent least-squares inverse. This procedure is a form of Newton’s method. The least-squares inverse of model sensitivity has been applied to many problems, e.g., as discussed by Bennett (1992).

Because the model is non-linear, the linearized forward model (sensitivity matrix) is incorrect when the guessed initial state is incorrect due to advection by an incorrect ‘background’ flow. These forward problem errors add to the other inadequacies of the forward model, but their effects can be assuaged by adding extra variance to the residual covariance R . As the estimate of the initial state is improved, re-linearizing the model can lead to a better forward model, so the extra ‘non-linearity’ errors can be gradually decreased when the inverse is repeated. The fixed-point iteration of ‘total inversion’

(Tarantola, 1987) is not implemented here, so the iteration proceeds relative to the last estimate, and the matrix R controls the step size. Large R means small steps toward the final solution, but taking too many steps will likely lead to overfitting and excess structure in the estimate. The linearized iteration is not guaranteed to converge, and a very poor starting guess for the model initialization can mean that the linearized forward model will drive the iteration in the wrong direction. Fortunately, the larger-scale structures tend to be more linear than the small-scale structures, and the dataset used here was adequate to allow the iteration to significantly improve the fit to observations in all cases.

At each step, the linear relation is inverted by choosing a statistically best solution in the least-squares sense.

$$\hat{\mathbf{m}} = PH^T(HPH^T + R)^{-1} \mathbf{d} = (H^T R^{-1} H + P^{-1})^{-1} R^{-1} \mathbf{d}. \quad (5)$$

At the end of each step, the non-linearity of the forward problem is checked by comparing the data values from a model run using the new initialization estimates to the linearized approximation used for the inverse at that step. For example, the ‘non-linearity’ for iteration i is the difference of the data from the starting and ending model runs in step i minus the data changes predicted by the linearization:

$$\mathbf{NL} = \mathbf{F}(\Psi_{i-1} + \Psi_i) - \mathbf{F}(\Psi_{i-1}) - H_i \hat{\mathbf{m}}_i \quad (6)$$

where Ψ_i is the state vector perturbation made from the basis function perturbations $\hat{\mathbf{m}}_i$ and Ψ_{i-1} is the guess at the start of the step, and was used to generate H_i , the linearized forward problem for the step.

Iteration was found to proceed most reliably if small steps were taken. We enforced this by increased $\underline{\sigma}$ so that the largest non-linearities were small compared to the assumed uncertainty. If the non-linearities were too large, the solution was discarded, and the inverse was repeated with larger uncertainty bounds. The iteration proceeds until either a fit is obtained within the expected uncertainty bounds, or the non-linearity and misfit cannot be reduced. If either the model or the forward problem is strongly non-linear (with respect to the uncertainty in the initial guess), then it may be necessary to abandon the iterated linearization and use a global grid search or Monte Carlo method, in order to avoid local minima in the least-squares fit. This was never attempted in our cases.

The inverse procedure also estimates uncertainty \hat{P} for the output model parameter values:

$$\hat{P} = P - PH^T(HPH^T + R)^{-1} HP = (H^T R^{-1} H + P^{-1})^{-1} \quad (7)$$

In the open ocean with no natural boundaries nearby, the boundary conditions should be measured and supplied at each timestep. In our situation, one can either use the unknown boundary conditions as adjustable parameters in the model, or extend the model domain so that the boundary conditions have little effect during the limited runtime needed to cover the data (Yakimew and Robert, 1990). The latter choice, which we adopt here (Fig. 3), transfers the uncertainty from the boundary conditions to the initial conditions. This is inefficient in some respects, but also provides a simple representation of features propagating into and out of the data domain.

To the extent that the model is non-linear, the dynamical evolution is incorrect in regions where the field is poorly determined, and is incorrect globally until the field is determined perfectly. For this reason, it is difficult to distinguish between model errors due to incorrect dynamics and due to poorly determined advecting fields. The error estimates from the least-squares procedure and the forecast provide linearized guesses for the error level expected if the dynamics were perfect and the prior statistical assumptions were correct (linear and Gaussian).

5. Identical twin experiments

As a test of the initialization by fitting, we conducted ‘identical twin’ experiments in the IFF model domain. In this framework we are able to check several issues. Is our selected set of basis functions adequate to correct a large fraction of the variance of an initial condition error field? Does the model fit converge towards the true initial state? Does a fit yield increased forecast skill?

We commenced a base run, the synthetic ‘observed’ data, from the initial time-independent OA of the non-synoptic Initialization Survey of the observed data. The velocity field was initialized with a 350-m level of no motion. It was then allowed to evolve with fixed temperature and salinity for three days before the ‘truth’ base run commenced (where this diagnostic spin-up was done solely for this twin case ‘observed’ run and no other). This ‘truth’ run was then sampled along the same cruise tracks as the Initialization Survey. From this synthetic non-synoptic ‘data’, we constructed a time-independent OA from which to commence our sensitivity runs. Note that since the model only has five vertical surfaces where temperature and salinity is defined, we chose to objectively analyze the T and S fields on the model surfaces themselves. This circumvents problems in defining vertical basis functions that map the model surfaces to three-dimensional fields. Additionally, we merged the synthetic OA initial state in the data region with the ‘true’ far-field initial conditions so that the error field of the initial conditions was confined to the data region. This circumvents the problem of supplying a climatology in the unknown far-field as must be done in reality.

5.1. Basis function efficiency

Since we know the initial condition error in the identical twin scenario, we evaluated the efficiency of the basis functions in explaining this error for each of the five model variables. Table 1 shows the variance explained by the basis function expansion of the initial condition error due to the time-independent OA of the cruise-track synthetic data. For scalar variables, 80–95% of the error variance is captured by 50 basis functions and 93–97% is captured by 100 functions. For the velocity components, 58–76% of the error variance is explained by 50 functions and 72–81% of the error variance is explained by 100 functions. The poorer performance of the basis for velocity is expected since the basis functions were derived for temperature. These results suggest that our choice of basis sets is likely to be an adequate descriptor for representing the true (unknown) initial condition error in the observations.

Table 1
Percent variance explained by structure functions

Identical twin case initial error fields		
Number of functions	50	100
<i>T</i> -1	85	93
<i>T</i> -2	82	93
<i>T</i> -3	85	93
<i>T</i> -4	94	96
<i>T</i> -5	95	97
<i>S</i> -1	86	94
<i>S</i> -2	85	94
<i>S</i> -3	86	94
<i>S</i> -4	95	97
<i>S</i> -5	95	97
<i>U</i> -1	71	81
<i>U</i> -2	64	81
<i>U</i> -3	70	79
<i>U</i> -4	67	72
<i>V</i> -1	67	80
<i>V</i> -2	58	78
<i>V</i> -3	68	81
<i>V</i> -4	76	80
<i>P</i> -1	93	96

T-*n* indicates temperature at model level *n*, *S*-*n* salinity at level *n*, *U*-*n*, *V*-*n*, baroclinic velocity components for mode *n*, and *P*-1 barotropic streamfunction.

5.2. Idealized initial condition error

We first tested the inverse strategy in several simple scenarios (e.g., a single basis function added to the ‘true’ initial condition in the temperature field of one layer). In this ideal case (initial condition error spanned completely by one basis function), the only thing preventing a perfect inverse solution is non-linearity. We were able to solve nearly exactly with one iteration for the amplitude of the basis function in perturbation runs involving one or a few basis functions as long as the perturbation was not so large ($< 1^\circ\text{C}$) as to induce strong non-linearity in response. In the case of significant non-linearity, several iterations were needed to converge to the true solutions in these simple cases.

We secondly constructed an initial condition error field as a 50-basis function truncation of the ‘true’ initial condition error for each variable in each layer. (These represent typically 80% of the ‘true’ initial condition error variance in most layers and for most variables as shown in Table 1.) When the model initial conditions are corrected by the inverse solution (1900 total unknowns), the model–data misfit variance is reduced to 87%. Since the linearized inverse predicts a 94% reduction in model–data misfit, the non-linearities in the forward problem are small. A comparison of the amplitudes of the basis functions predicted by the inverse vs. the actual amplitudes used in the construction of the initial difference fields shows that the larger amplitude coefficients usually agree within a factor of two while the smaller amplitude ones can

differ appreciably. Therefore, for each day of the model run, we evaluated the ‘model–data field error variance’ defined to be the instantaneous squared error (‘truth’ run minus adjusted initial condition run) averaged over all grid points in the data domain (12°W–9.5°W; 63.5°N–65°N). The temperature and salinity field error variance was typically reduced by 75% throughout the hindcast and also several days into the time of independent data. This shows that the inverse is converging towards the true solution and not just fitting the data. The iteration did not converge completely, however, but found a local minimum in the objective function due to indeterminacy; a non-linear search may have been needed to find the exact solution. The velocity field error variance was only reduced 20–70% in the hindcast and forecast interval suggesting that the hydrographic data is more efficient at constraining hydrography rather than velocity.

5.3. Full initial condition error

We also tested the inverse method using the raw initial condition error, so the set of basis functions spans only a part (80–95%, Table 1) of the complete initial condition error variance. We computed the sensitivity of the synthetic-OA initial state to each of 100 basis functions for all model variables (1900 total unknowns). When the model initial conditions are corrected by the inverse solution, the model–data misfit variance is reduced 85%, compared to the expected reduction of 91% predicted by the linearized inverse. This improvement in fit is approximately the same as for the truncated (and completely spanned) initial condition error case (Section 5.2) and supports the idea that basis function completeness is not the limiting factor in fitting the data.

Fig. 4 shows the reduction in along-track model–data temperature misfit to illustrate the performance of the inverse in fitting the model to the data as a function of time and layer. The major differences between the synthetic ‘observed’ and modeled measurements are near the frontal crossings. The size of the differences between the dotted and thick black line is an indication of the strength of non-linearity. We chose sufficiently large spatially constant elements for σ (3°C and 0.2 psu) to keep the step-wise-non-linearity small (Section 4).

We again tested for convergence towards the true solution by computing the reduction in model–data field error variance (defined in Section 5.2). Fig. 5 shows the field error variance for temperature as a function of day and layer for the base run (OA initial conditions), persistence of the OA and the inverse solution iteration. The inverse solution is effective at reducing the field error variance by typically a factor of two or more in the fitting interval (days 0 to 3). The results are similar for salinity and velocity. This shows that for reasonable strength of the non-linearity the fitting procedure works and is converging towards the true solution. We iterated a second step and achieved an additional misfit variance reduction of 64% (compared to 68% predicted by the linear inverse). But as can be seen in Fig. 5 there was little reduction in field error variance. So even though the model–data fit was significantly improved, the model state estimate was not improved. This is similar to results in Section 5.2, where the initial condition error was completely spanned by the basis functions but improved data fit did not improve model state. This suggests that there is insufficient data to drive the estimate any closer to the true initial condition, although many statistical assumptions affect the fit.

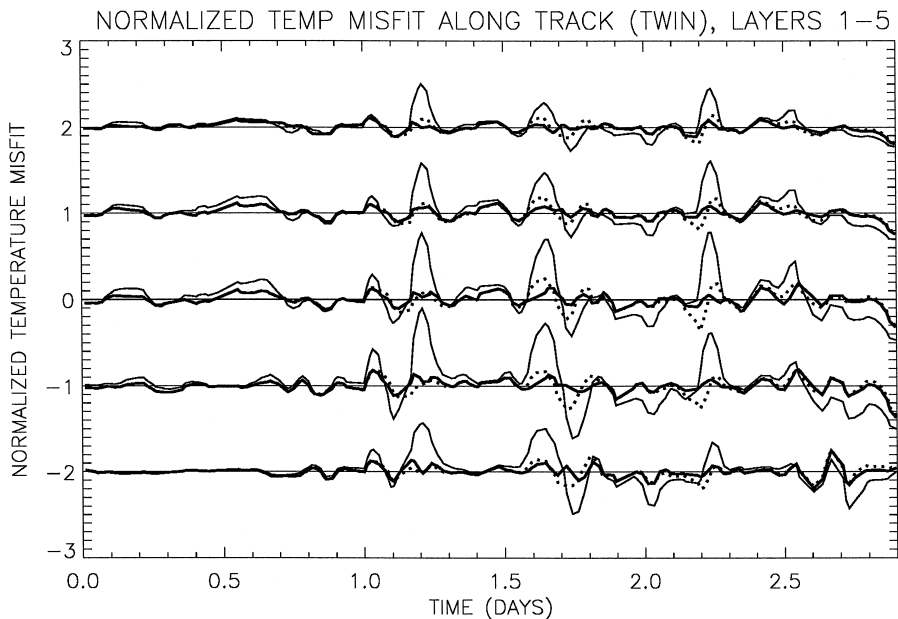


Fig. 4. Twin experiment: time series of normalized temperature model–data misfit (normalized by 3°C) following the cruise track of the Initialization Survey for (top to bottom) model layers 1 to 5 (each offset by a constant value). Thin line is model–data misfit when initializing run from a twin OA of the Initialization Survey synthetic hydrographic data. Thick line is misfit predicted from the linear inverse solution. Dotted line is actual model–data misfit after adjusting initial conditions. The differences between predicted and actual misfit indicate non-linearity in the model response and are concentrated near the frontal crossings.

5.4. Forecast skill

We next examined the skill levels of the forecast from adjusted initial conditions of Section 5.3 in the time interval of independent data. We use persistence of the Initialization Survey OA as the baseline skill to beat. As shown in Fig. 5, the model forecast generally beats persistence on day four and five of the forecast, corresponding to the first legs of the zig-zag forecast. By day 6, however, the model run does not beat persistence in any layer. Longer term forecasts also did not exhibit any appreciable skill.

If we use only the Zig-Zag Survey ‘data’ samples to verify the field forecast skill shown by the model for days 4 and 5, we find that the along-track measure of skill is not better than the persistence forecast even though the field is better. The sparsely sampled data of the zig-zag track is inadequate to show the forecast skill. Obviously, the one realization studied here cannot produce a verdict about the usefulness of limited verification data. But the one forecast result shown here using ‘identical twin’ experiments indicates that the limited non-synoptic validation data in this rapidly evolving frontal realization may be inadequate for revealing forecast ‘skill’ (at least by our measure). Perhaps satellite data or rapid aerial XBT surveys which can provide more complete spatial coverage in short time intervals could do better in future validation studies.

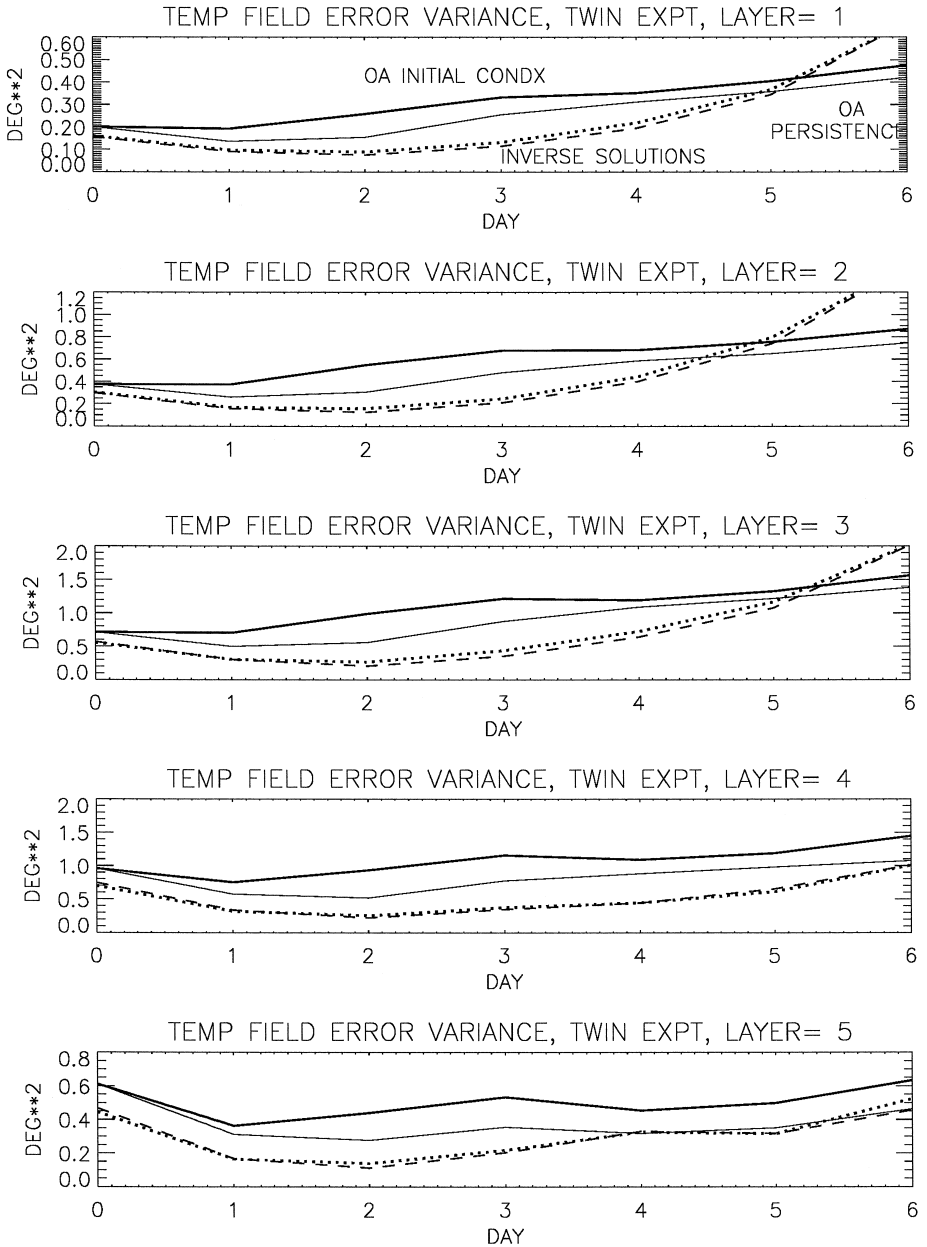


Fig. 5. Twin experiment: temperature ‘field error variance’ (Section 5.2) in the entire data domain for (thin line) persistence of the Initialization Survey OA, (thick line) model run initialized from the Initialization Survey OA, (dotted line) first iteration model run initialized with the inverse solution and (dashed line) second iteration model run initialized with the inverse solution for (top to bottom) layers 1–5. Note the differing scales on the abscissas. The time interval from 0–3 days corresponds to the fitting interval. The time interval from 3–6 days corresponds to forecasting independent data.

6. IFF August 1993 dataset experiments

The three hydrographic surveys taken in the IFF in August 1993 provide us with several scenarios of hindcasting and forecasting:

Case 1: Fit model to Initialization Survey.

Case 2: Forecast Zig-Zag Survey from Case 1.

Case 3: Forecast Validation Survey from Case 1.

Case 4: Fit model to Initialization and Zig-Zag Survey.

Case 5: Forecast Validation Survey from Case 4.

Case 6: Fit model to Zig-Zag Survey.

Case 7: Forecast Validation Survey from Case 6.

Case 8: Fit model to Zig-Zag and Validation Surveys.

Cases 1, 4, 6, and 8 correspond to hindcasts of 3-day, 6-day, 2-day, and 6-day time intervals, respectively. These cases provide tests of the fitting procedure, the consistency of model physics with nature, and the strength of the non-linearity of the system. They also provide initial conditions for Cases 2, 3, 5, and 7 that correspond to forecasts (executed retrospectively) into data-independent intervals. Although verification data is non-synoptic, the four forecast cases correspond, respectively, to approximately 2–3 day forecast, a 7-day forecast, a 4-day forecast, and a 4-day forecast in the most energetic region of the IFF.

6.1. Case 1. Fitting the Initialization Survey

Our first experiment (Case 1) with the observations is to fit the model to the Initialization Survey. We computed the forward problem by perturbing the initial conditions from the OA of the non-synoptic survey. We executed one perturbed run for each of the 100 basis functions for each layer and each dependent variable (temperature, salinity, baroclinic zonal and meridional velocity, and barotropic streamfunction). We adjusted the diagonal of R (expected residual covariance) to control the step size to ensure that non-linearity is small in the linearized inverse. With rms expected residuals set to 3°C , the misfit variance of the linear inverse approximates to within 10% the misfit variance of the model run from the corrected initialization implying weak non-linearity. The non-linearity of the forward problem is evident in plots of along-track model–data misfit as the difference between the actual misfit from the corrected model run and the misfit predicted by the linearized inverse (Fig. 6). The non-linearity is largest at the frontal crossings. Larger values of the expected residual variance reduce the influence of the non-linearities on the step-wise solution.

In the first case, we not only tested spatially constant elements for σ but also spatially variable elements. We first tried to set the largest uncertainty at the initial frontal position from the OA. However, the front evolved too quickly in time for this to be successful. We secondly used the model run to estimate the uncertainty by calculating the non-linearity for each datum and assigning proportional uncertainty. This worked reasonably well and decreased the misfit variance by 5–7% over the case of uniform 3°C uncertainty by giving more weight to quiet regions far from the front. On the other hand, it is complicated to describe and did not give huge improvements so we did not use the technique in any other fit. We now discuss in more detail the results of this case.




Fig. 6. Time series of normalized temperature model–data misfit following the cruise track of the Initialization Survey for (top to bottom) model layers 1 to 5 (each offset by a constant value). The misfit normalization varies from 1°C away from the front to 3°C near the front. Thin line is misfit in model when initializing run from an OA of the Initialization Survey hydrographic data. Thick line is misfit predicted from linearized inverse solution. Dotted line is actual model–data misfit after adjusting initial conditions. The differences between predicted and actual misfit indicate non-linearity in the model response and are concentrated near the frontal crossings.

The first step used spatially varying σ , roughly 3°C and 0.1 psu near the front and 1°C and 0.1 psu in the far-field. The resulting corrected initial state reduced the normalized model–data misfit by 65% (Fig. 7). This is close to the misfit variance reduction of 71% predicted by the inverse. The largest remaining misfits are centered around the times of the frontal crossings as would be expected for a non-linearly evolving field. The fact that we can fit the model at this skill level in one iteration suggests non-linearity is not an insurmountable problem in the 3-day time scale of the hindcast. The results of the twin experiments suggest we are converging onto the true ocean state and that the model physics seems adequate to represent this flow.

The second iteration used spatially constant elements for σ (3°C and 0.1 psu) and reduced the remaining misfit variance by 33% compared to the 45% reduction predicted by the linear inverse. The total misfit variance reduction is then 77%, approaching the limits expected for the basis function expansion. Fitting the data more closely may lead to spurious structure in the initialization due to fitting environmental noise in the observations.

Additional iterations of the inverse method slightly improved the fit (Fig. 7) but, as we shall see next, did not improve the forecast skill, so the iteration procedure was terminated.

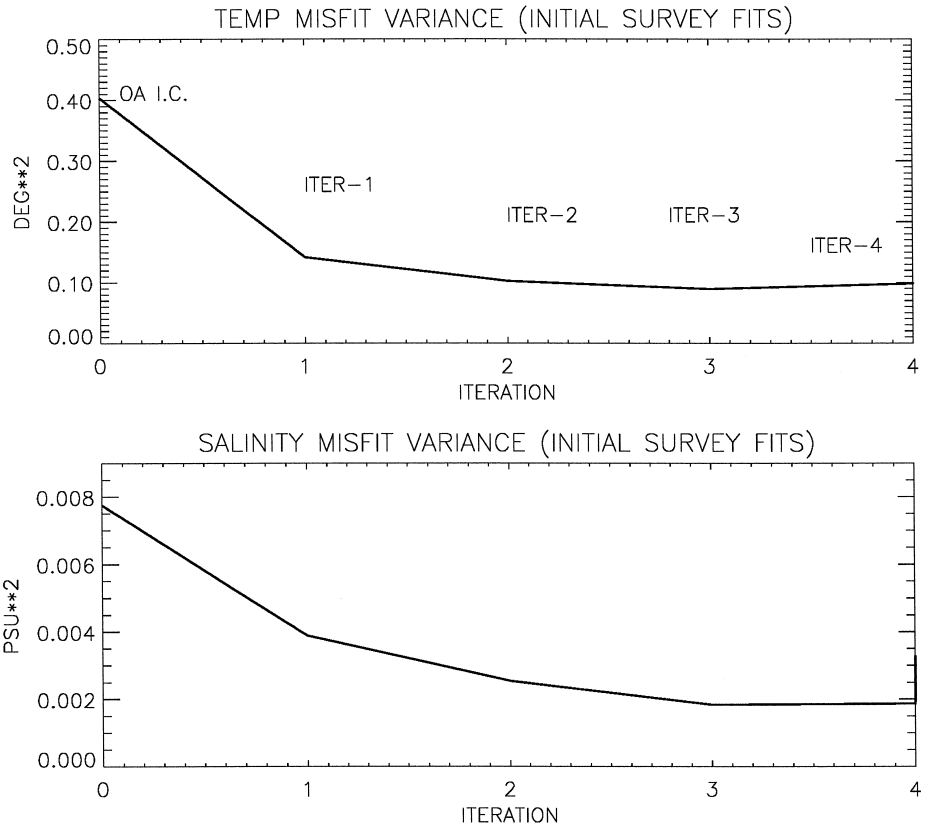



Fig. 7. (Top) Temperature and (bottom) salinity misfit variance of Case 1 fits of the Initialization Survey hydrographic data. The first point corresponds to persistence of the Initialization Survey OA. The second through fifth points correspond to the first through fourth iterations of the model fits.

6.2. Cases 2 and 3. Forecasting from Initialization Survey

We quantitatively evaluated the forecast skill (Case 2) of the model fit of the Initialization Survey by computing the rms temperature forecast error along the zig-zag track. The base run (OA, initial conditions) had an rms forecast error of 1.25°C, the first inverse had 1.55°C and the second inverse had 1.60°C. This drop in forecast skill is

Fig. 8. Time sequence of 75 m model temperature for August 14–19, corresponding to the day 0–5 of the model fits and forecasts. (Top row) Initial conditions from OA of Initialization Survey. (2nd row) First iteration of Case 1 inverse solution fit to Initialization Survey. (3rd row) Third iteration of Case 1 inverse solution fit to Initialization Survey. (4th row) Third iteration of Case 4 inverse solution fit to both Initialization and Zig-Zag Surveys. Contour intervals are 1°C. White is warmer than 7°C, light grey warmer than 5°C, dark grey warmer than 3°C and black colder than 1°C.



evident also in the qualitative behavior of the inverse solution forecast in that the model tends to produce a zonal or northeastward flow in the center of the data domain by August 18 and 19 (Fig. 8, second and third rows), yet the observed Zig-Zag Survey clearly shows a southeastward flow (Miller et al., 1995b; Robinson et al., 1996a). These results suggest to us that the model is able to fit the data but it is doing so at the expense of capturing the true nature of the evolution. Because of limited initialization data, there are multiple solutions that fit the observations. Adding additional data (e.g., the Zig-Zag Survey itself) can refine the choice of initialization and produce quantitative agreement and realistic qualitative behavior. As we shall see later, the model physics do allow a fit to both the Initialization and Zig-Zag Surveys, which argues against grossly incorrect model physics or inadequate basis function representation.

The forecast from the OA initial conditions appeared to be more qualitatively correct than the forecasts from the inverse solutions. This is consistent with the results of Miller et al. (1995b) who showed that quasigeostrophic forecasts launched from non-synoptic OA initial conditions had quantitative skill if the forecast validation time is computed based on the local time difference between the non-synoptic initialization and validation data. Indeed, the PE model run here initialized from the OA produced the southeastward current 2 days after the run commenced. Thus, if one simply ascribes August 16 (the day data was collected in the western third of the domain) to the initial conditions, the 2-day forecast for August 18 would qualitatively match the observed data (Fig. 8, top row). Even if our inverse results had exhibited qualitative forecast skill, we would not be surprised to find reduced quantitative forecast skill during the Zig-Zag Survey because the twin experiments showed it to inadequately sample the rapidly evolving frontal structure. A modified design for the non-synoptic Zig-Zag survey or other synoptic datasets (SST, altimetry, aerial XBT surveys) might provide better assessments of quantitative skill measures.

Since the model fits had no apparent skill in the Zig-Zag Survey, there is no reason to anticipate that the model has skill in the Verification Survey (Case 3) which is indeed the case here.

6.3. Case 4. Fitting the initialization and Zig-Zag Surveys

We next attempt to fit the model to both the Initialization and the Zig-Zag Surveys, a time interval spanning 6 days. Rather than starting from OA initial conditions, we began the first set of perturbation runs from the first iteration of Case 1 (the fit to the Initialization Survey alone). We ran 100 basis functions for each dependent variable in each layer as for Case 1. We chose to use spatially constant elements for σ for temperature (4°) and salinity (0.3 psu) for simplicity.

Fig. 9 shows the decrease in temperature and salinity misfit variance for each of the three iterations we performed, along with the results from OA initial conditions of Case 1 and the first iteration of Case 1 (from which we launched this set of iterations). The temperature misfit variance decreases with each iteration, while the salinity misfit variance decreases with the first iteration but does not decrease much for the second two iterations. After three iterations, the combined temperature and salinity misfit variance is reduced 67% compared to the initial hindcast. This suggests the model physics are

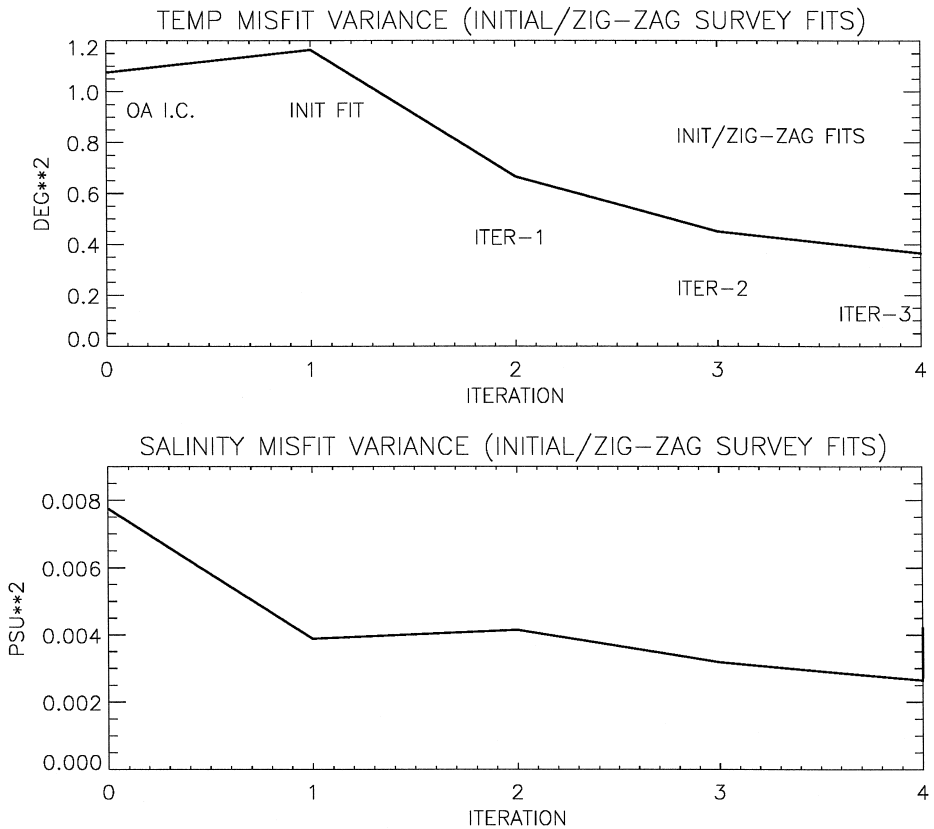
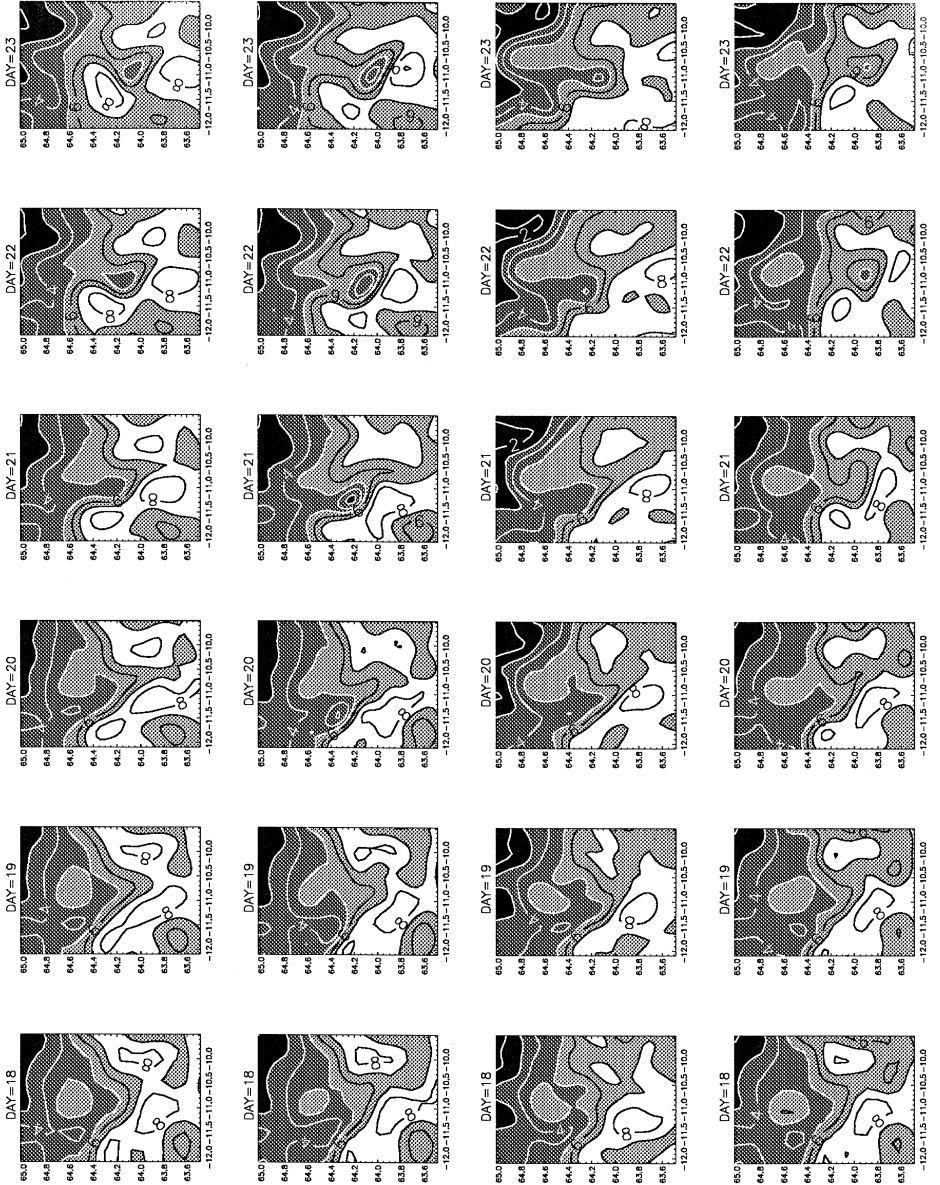


Fig. 9. (Top) Temperature and (bottom) salinity misfit variance of fits of the Initialization and Zig-Zag Surveys hydrographic data. The first point corresponds to initial conditions from the Initialization Survey OA. The second point corresponds to the first iteration of Case 1 inverse solution fit to only Initialization Survey data. The third through fifth points correspond to the first through third iterations of the Case 4 model fits to both surveys.

capable of representing the flow conditions in a 6-day period of a vigorous current instability.

Because of the limited observations, mere quantitative fits might mean that the model could converge to a dynamically incorrect initial state. This does not seem to be the case, since the time evolution of the flow (Fig. 10) compares well visually with the expected structures seen from the individual OA's of the two surveys (Fig. 2). The frontal current meanders southeastward and an isolated eddy to the north of the front moves northeastward during the 6-day time interval. This eddy was apparent in the initial OA but its spatial structure was not captured by the survey. The model has chosen the proper spatial scale to allow the eddy to migrate in a fashion that is dynamically consistent with what was observed. These points suggest that the model is indeed converging to a physically reasonable state, capturing the physical evolution of the frontal system in the 6-day interval.



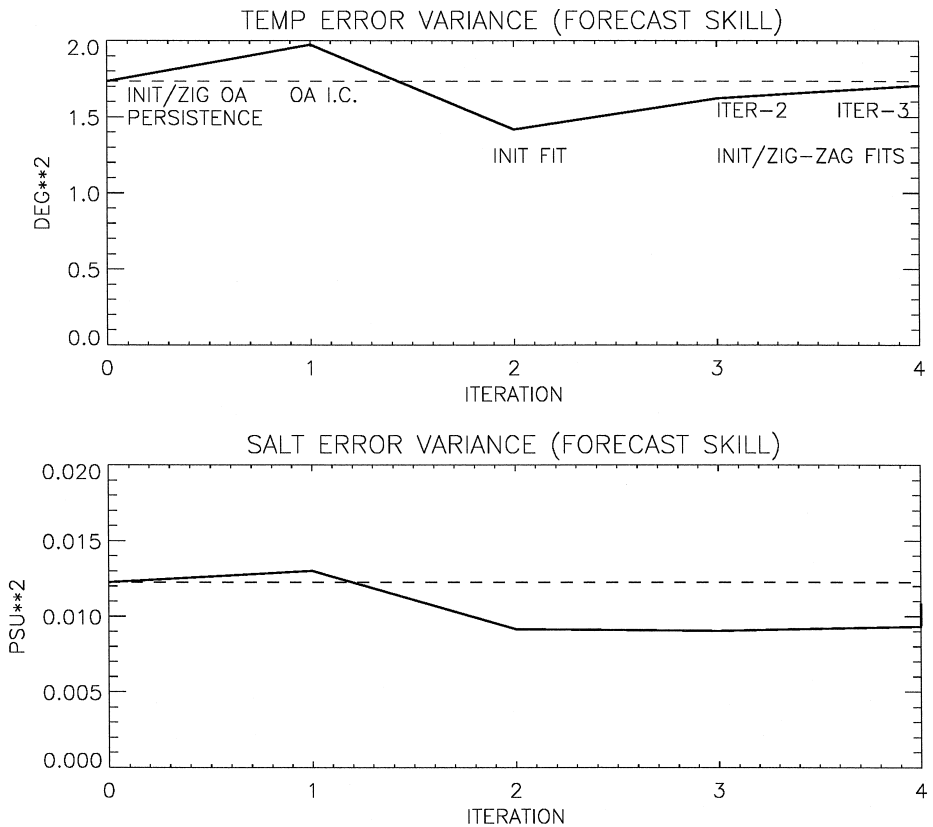


Fig. 11. (Top) Temperature and (bottom) salinity forecast error variance for the Validation Survey hydrographic data. The first point corresponds to persistence of the Zig-Zag/Initialization Surveys OA and is indicated by the dashed line. The second point corresponds to the model forecast initialized by the OA of the first point. The third point is the forecast from the Case 1 fit of only the Initialization Survey; this forecast, although it has the best 'skill', does not represent the hammerhead instability at all. The last two points represent the second and third iterations of the Case 4 fits of the Initialization and Zig-Zag Survey data; the second iteration qualitatively captures the hammerhead instability, while the first and third iterations do not.

6.4. Case 5. Forecasting from Zig-Zag Survey

We next discuss the skill in forecasting the Validation Survey from the model fits to the first two surveys. Fig. 11 shows the temperature and salinity forecast error variance associated with the forecasts of several of these fits. For a baseline comparison, we show

Fig. 10. Time sequence of 75 m model temperature for August 18–23. (Top row) Initial conditions from OA of Zig-Zag/Initialization Surveys blend. (2nd row) Third iteration of inverse solution Case 6 fit to Zig-Zag Survey. (3rd row) Third iteration of inverse solution Case 4 fit to Initialization and Zig-Zag Surveys. (4th row) Fourth iteration of inverse solution Case 8 fit to both Zig-Zag and Validation Surveys. Contour intervals are 1°C. White is warmer than 7°C, light grey warmer than 5°C, dark grey warmer than 3°C and black colder than 1°C.

the skill of the forecasts of persistence of the initial OA (where the OA includes the Zig-Zag, weighted more strongly, and the Initialization Survey). The typical model forecast skill is not very different from the persistence forecast. The two points on the right of Fig. 11 are the forecast skill (Case 3) of the second and third iteration of Case 2. The two forecasts exhibit temperature forecast skill slightly better than persistence, and salinity forecast skill that is only slightly better than the temperature results. The forecast initialized from the Zig-Zag/Initialization OA (second point on Fig. 11) has the worst skill in this scenario but that measure might be altered if a different initial time is specified for the non-synoptic OA (an effect we did not explore).

It must be emphasized, however, that the forecast initialized from Case 1 (center point of Fig. 11) has the highest data forecast skill in spite of the fact that the model run (not shown) does not resemble the structures seen in the time sequence of OAs. Thus, a poor initial state can accidentally generate quantitative forecast skill that is misleading. This is likely a consequence of computing skill statistics on a front, so a smooth forecast (as in Case 1) with a frontal slope may match an observed frontal shape with higher fidelity than a forecast which has a lot of eddy structure around the front. Thus, we cannot ascribe much significance to this particular quantitative skill measure in our study, although other measures of skill might yield more favorable results.

There is clearly, however, qualitative skill in the two-survey model fits that gives substance to the quantitative verification. The hammerhead meander of the frontal current found by Miller et al. (1995b) and Robinson et al. (1996a) is also found in our forecasts here. The eddy north of the front that migrates eastward and merges with the frontal current is also captured by the simulation (Fig. 10). Since the original forecasts of Robinson et al. (1996a) were not available for comparison with these results we cannot state which forecasts have more skill and our confidence in such a differentiation would be small.

6.5. Case 6. Fitting the Zig-Zag Survey

We also fit the model to the Zig-Zag Survey alone. This sequence commenced from the Zig-Zag/Initialization OA. We used relatively small spatially constant elements for σ (2°C and 0.1 psu) since the adjusted initial condition runs were very linear. The misfit variance for the first iteration (50 basis functions) was reduced 64% compared to the predicted reduction of 69%. This shows the model is behaving linearly in the this 2-day time interval. The second and third iterations (100 basis functions each) reduced the remaining misfit variance by 29 and 30%, respectively (with 31% predicted misfit variance reduction each iteration). The combined total misfit variance reduction for the three iterations is then 82%. The 2-day time sequence for each iteration retains observationally believable structures (Fig. 10, second row).

6.6. Case 7. Forecasting from the Zig-Zag Survey

The structures that develop after the 2-day model fit to the Zig-Zag Survey exhibit the hammerhead instability described by Miller et al. (1995b) and Robinson et al. (1996a). Since the 6-day fits using the Initialization and the Zig-Zag Surveys failed to

consistently forecast a hammerhead instability, we conclude that it was crucial to have the precise fit to the Zig-Zag Survey to obtain the observed instability. This sensitive dependence on initial conditions is consistent with the known baroclinically unstable dynamics of the IFF (Miller et al., 1995a,b).

In spite of their appealing qualitative skill, the Case 7 forecasts of the Validation Survey are unable to beat persistence of the Zig-Zag/Initialization OA. Again it should be emphasized that there may be inadequate validating data to properly assess the true skill of the forecast and because the strong temperature gradients of the frontal region can mask differences in small-scale but energetically important structures.

6.7. Case 8. Fitting the Zig-Zag and Verification Surveys

We lastly explore the capability of the model to fit the Zig-Zag and Verification Surveys. The development of the hammerhead (or deep-sock) intrusion is a non-linear

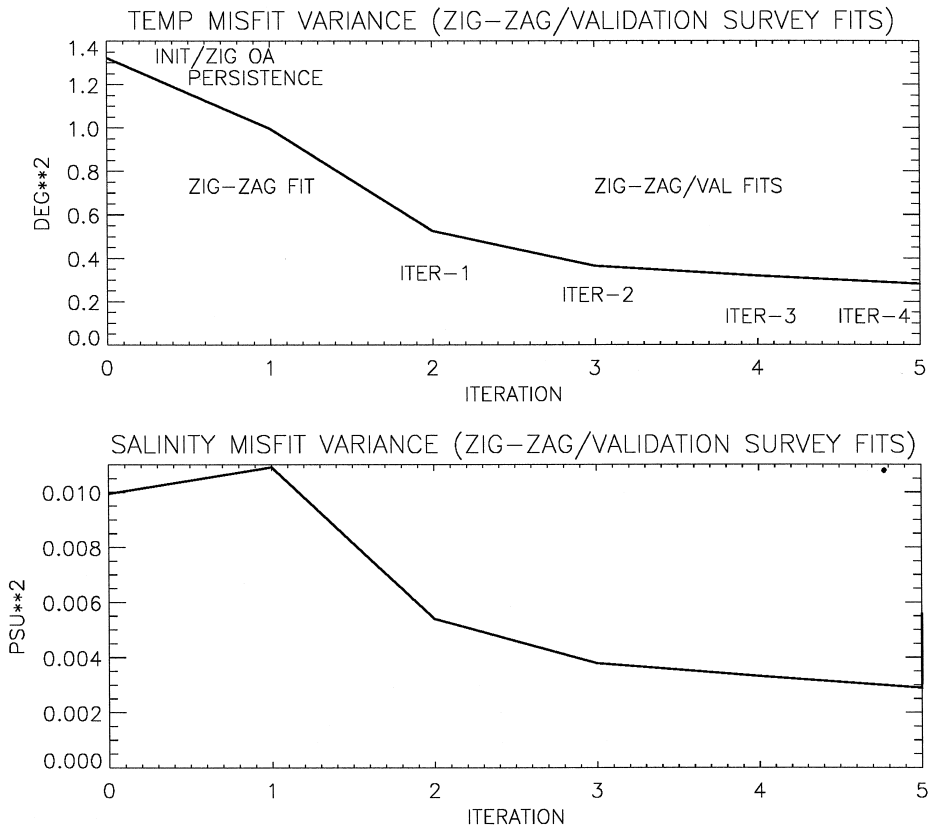


Fig. 12. (Top) Temperature and (bottom) salinity misfit variance of fits of the Zig-Zag and Validation Survey hydrographic data. The first point corresponds to persistence of the Zig-Zag/Initialization Surveys OA. The second point corresponds to the second iteration of the Case 6 model fit using the zig-zag data alone. The third through sixth points correspond to the first through fourth iterations of the Case 8 model fits.

baroclinic instability process (Miller et al., 1995b) which was forecast in real time by Robinson et al. (1996a). However, the quantitative forecast skill shown by Robinson et al. (1996a) was limited to the deeper layers of the models because the observations had a baroclinic tilt to the instability that the model failed to capture.

Fig. 12 shows the temperature and salinity misfit variance for each iteration of the fits to the Zig-Zag and Validation Surveys. The major reduction of temperature and salinity misfit variance occurs in the first iteration (49% relative to the first inverse solution for the Zig-Zag Survey alone). By the fourth iteration the total misfit variance is reduced by 74%.

The spatial structure of the best fit to the two surveys is shown in Fig. 10 (bottom row). Its evolution resembles the results of Robinson et al. (1996a) except here the amplitude of the instability is smaller and more realistic. The eddy north of the front, which was absent from the forecast of Robinson et al. (1996a), is well-modeled and retains its structure north of the hammerhead as seen in the satellite SST image of Robinson et al. (1996a). On the other hand, the baroclinic shear of the hammerhead, evidenced by the 25 m temperature evolving slightly upstream of the deeper temperatures, was not captured by this hindcast or by the forecast of Robinson et al. (1996a,b), suggesting the inability of the coarse vertical resolution and/or the model physics to simulate this effect.

7. Summary and discussion

A fitting procedure was tested and applied with the Harvard PE model to non-synoptic hydrographic surveys of unstable current meandering of the IFF. The initial conditions (including the fields outside the data domain) were adjusted with eddy-scale basis functions to optimize the model fit to the observations, with no additional forcing or adjustment of boundary conditions during the model runs.

The technique was first tested with an ‘identical twin’ predictability experiment. This showed the inverse technique can successfully fit the non-synoptic Initialization Survey data, correct a large fraction of the initial condition error, and allow the model to move closer to the true evolution. However, although additional iterations of the fitting procedure could improve the fit to the Initialization Survey, the model could not be adjusted closer to the true initial state because of the limited initialization data. Moreover, the limited verification data was inadequate to show unambiguous forecast skill even for short 2-day forecasts which were known to have skill in the identical twin framework.

The PE model was then fit to the observed hydrographic data from August 1993 in several scenarios. With only a few iterations, temperature and salinity model–data misfit variance was reduced 70–80% relative to initializing the model from a time-independent OA. The success of the fit was supported by qualitative realism of the frontal variability as described previously (Miller et al., 1995b; Robinson et al., 1996a,b).

The model run from the optimized initialization constitutes a possible dynamically consistent scenario explaining some of the variability seen in the IFF observations as a meandering of the front. To the extent that the model is accurate, the observations have

been reconstructed into a four-dimensional picture of the flow field in the area. The results here set the stage for diagnostic analyses of the frontal baroclinic instabilities, e.g., using a PE analogue (presently under development by P. Lermusiaux, A. Miller, N. Pinardi and A. Robinson) of the energy and vorticity analysis (EVA) system developed by Pinardi and Robinson (1986) for quasigeostrophic dynamics and applied by Miller et al. (1995a,b) to quasigeostrophically modeled IFF variations.

Although hindcast skill increased, quantitative forecast skill (measured by error variance) was not always increased relative to the time-independent OA initialization. Qualitative skill assessment proved necessary to distinguish the integrity of the hindcasts and forecasts, especially the occurrence of a hammerhead baroclinic instability of the IFF (Miller et al., 1995b; Robinson et al., 1996a). Since the model was able to be successfully fit to the hammerhead instability, the incorrect forecasts of the hammerhead from the Zig-Zag Survey fits are most likely a consequence of inadequate initialization data. The highly non-linear IFF variability leads to major differences in the evolution of flow from slightly different initial states (Fig. 10). This suggests that finely resolved and nearly synoptic hydrography is necessary in the IFF to constrain the initialization. The successful hindcasts argue against inadequate model physics or incomplete basis functions as being the major factors limiting forecast skill. The quantitative forecast skill shown by Robinson et al. (1996a), who used an optimal interpolation technique in real time, is comparable to what was obtained here for the hammerhead when the model was initialized from or fit to the antecedent (Zig-Zag) survey.

The main difference between this method and the representor method of Bennett (1992) is our use of the truncated set of basis functions. An adjoint model can give the grid point structure of the data sensitivity in the forward problem, but using smoothing assumptions (i.e., a covariance matrix) similar to those used here should lead to similar solutions. Due to non-linearity, the small-scale structure of the sensitivity is often less reliable than the large-scale structure. This is a justification for only fitting the larger scale structures.

An ocean model is a valuable tool for the interpolation and interpretation of data, as well as practical prediction problems, but it is dependent on obtaining adequate data for judging model quality. The techniques tested here could easily accommodate other data types, but the best type of data for constraining the model initializations and for testing forecasts skill would appear to be finely resolved synoptic hydrographic surveys as could be obtained from aerial XBT surveys. For this particular IFF dataset, there are also drifter observations, current meter observations, and a satellite SST image which can be used in future applications of this technique to improve the fits and forecasts. Higher vertical resolution in the model is of greatest priority in improving the details of the model fit, especially the baroclinic structure of the hammerhead instability.

Acknowledgements

This work was funded by ONR (N0014-96-0264) and NOAA (NA36GPO372 and NA47GPO188). We thank Prof. Allan Robinson for allowing our use of the HOPS PE model and Pierre Lermusiaux, Hernan Arango, Wayne Leslie, Pat Haley and Carlos

Lozano for aiding in our application of the model. The detailed and incisive comments of the anonymous reviewer, Phil Bogden (as a reviewer) and Pierre Lermusiaux greatly improved the presentation and interpretation of results. AJM thanks Prof. Allan Robinson for the many lively and provocative discussions on ocean physics and forecasting that were initiated during AJM's days at SACLANTCEN and that have continued thereafter in many colorful locales.

References

- Allen, J.T., Smeed, D.A., Chadwick, A.L., 1994. Eddies and mixing at the Iceland–Faeroe Front. *Deep-Sea Res.* 41, 51–79.
- Bennett, A.F., 1992. *Inverse Methods in Physical Oceanography*, Monographs on Mechanics and Applied Mathematics. Cambridge Univ. Press, New York, 346 pp.
- Bogden, P.S., Davis, R.E., Salmon, R., 1993. The North Atlantic circulation—combining simplified dynamics with hydrographic data. *J. Mar. Res.* 51, 1–52.
- Bogden, P.S., Malanotte-Rizzoli, P., Signell, R., 1996. Open-ocean boundary conditions from interior data—local and remote forcing of Massachusetts Bay. *J. Geophys. Res.* 101, 6487–6500.
- Bretherton, F.P., Davis, R.E., Fandry, C.B., 1976. A technique for objective analysis and design of oceanographic experiments applied to MODE-73. *Deep-Sea Res.* 23, 559–582.
- Bryan, K., Cox, M.D., 1967. A numerical investigation of the oceanic general circulation. *Tellus* 19, 54–80.
- Ghil, M., Malanotte-Rizzoli, P., 1991. Data assimilation in meteorology and oceanography. *Advances in Geophysics* 33, 141–266.
- Hopkins, T.S., 1991. The GIN Sea. Review of physical oceanography and literature from 1972. *Earth Sci. Rev.* 30, 175–318.
- Le Dimet, F.-X., Talagrand, O., 1986. Variational methods for analysis and assimilation in meteorological observations. *Tellus A* 38, 97–110.
- Lermusiaux, P.F.J., 1997. *Error Subspace Data Assimilation Methods for Ocean Field Estimation: Theory, Validation and Applications*. PhD dissertation, Harvard University, Cambridge, MA, 402 pp.
- Lermusiaux, P.F.J., Anderson, D.G.M., Lozano, C.J., 1998. Error and variability subspace initial estimates and the mapping of multivariate geophysical fields. *Q. J. R. Meteorol. Soc.*, sub judice.
- Lozano, C.J., Haley, P.J., Arango, H.G., Sloan, N.Q., Robinson, A.R., 1994. *Harvard coastal/deep water primitive equation model*. Reports in Meteorology and Oceanography: Harvard Open Ocean Model Reports, 15 pp.
- Miller, A.J., Arango, H.G., Robinson, A.R., Leslie, W.G., Poulain, P.-M., Warn-Varnas, A., 1995a. Quasi-geostrophic forecasting and physical processes of Iceland–Faeroe Frontal variability. *J. Phys. Oceanogr.* 25, 1273–1295.
- Miller, A.J., Poulain, P.-M., Robinson, A.R., Arango, Leslie, W.G., Warn-Varnas, A., 1995b. Quantitative skill of quasigeostrophic forecasts of a baroclinically unstable Iceland–Faeroe Front. *J. Geophys. Res.* 100, 10833–10849.
- Miller, A.J., Lermusiaux, P.F.G., Poulain, P.-M., 1996. A topographic-Rossby mode resonance over the Iceland–Faeroe Ridge. *J. Phys. Oceanogr.* 26, 2735–2747.
- Niiler, P.P., Piasek, S., Neuberger, L., Warn-Varnas, A., 1992. Sea-surface temperature variability of the Iceland–Faeroe Front. *J. Geophys. Res.* 97, 17777–17785.
- Orlanski, I., 1976. A simple boundary condition for unbounded hyperbolic flows. *J. Comput. Phys.* 41, 251–269.
- Perkins, H., 1992. *Large-scale structure of the Iceland–Faeroe Front*. SACLANT Undersea Research Centre Report SR-189, La Spezia, Italy, 40 pp.
- Pinardi, N., Robinson, A.R., 1986. Quasigeostrophic energetics of open ocean regions. *Dyn. Atmos. Oceans* 10, 185–219.
- Poulain, P.-M., Warn-Varnas, A., Niiler, P.P., 1996. Near surface circulation of the Nordic Seas as measured by Lagrangian drifters. *J. Geophys. Res.* 101, 18237–18258.

- Robinson, A.R., 1996. Physical processes, field estimation and an approach to interdisciplinary ocean modeling. *Earth-Sci. Rev.* 40, 3–54.
- Robinson, A.R., Arango, H.G., Miller, A.J., Poulain, P.-M., Warn-Varnas, A., Leslie, W.G., 1996a. Real time operational forecasting on shipboard of the Iceland–Faeroe frontal variability. *Bull. Am. Meteorol. Soc.* 77, 243–259.
- Robinson, A.R., Arango, H.G., Warn-Varnas, A., Leslie, W., Miller, A.J., Haley, P.J., Lozano, C.J., 1996b. Real-time regional forecasting. In: P. Malanotte-Rizzoli (Ed.), *Modern Approaches to Data Assimilation in Ocean Modeling*. Elsevier, New York, pp. 377–410.
- Robinson, A.R., Lermusiaux, P.F.G., Sloan, N.Q., III, 1998. Data assimilation. In: Brink, K.H., Robinson, A.R. (Eds.), *The Sea*, Vol. 10. Wiley, New York, pp. 541–594.
- Scheinbaum, J., Anderson, D.L.T., 1990. Variational assimilation of XBT data: Part I. *J. Phys. Oceanogr.* 20, 672–688.
- Semtner, A.J., 1974. An oceanic general circulation model with bottom topography. Tech Rep. 9, Department of Meteorology, University of California, Los Angeles, 99 pp.
- Spall, M.A., Robinson, A.R., 1989. A new open ocean, hybrid coordinate primitive equation model. *Math. Comput. Simul.* 31, 241–269.
- Tarantola, A., 1987. *Inverse Problem Theory: Methods for Data Fitting and Model Parameter Estimation*. Elsevier, New York.
- Tokmakian, R., 1994. The Iceland–Faeroe Front: a synergistic study of hydrography and altimetry. *J. Phys. Oceanogr.* 24, 2245–2262.
- van Leeuwen, P.J., Evensen, G., 1996. Data assimilation and inverse methods in terms of a probabilistic formulation. *Mon. Weather Rev.* 124, 2898–2913.
- Wunsch, C., 1996. *The Ocean Circulation Inverse Problem*. Cambridge Univ. Press, 442 pp.
- Yakimew, E., Robert, A., 1990. Validation experiments for a nested grid-point regional forecast model. *Atmos.-Ocean* 28, 467–472.