Version of Record: https://www.sciencedirect.com/science/article/pii/S027311771930660X
Manuscript_efb72e8d9f0c60f5304d088697a37e10

The Impact of Remote Sensing Observations on Cross-Shelf Transport Estimates from 4D-Var Analyses of the Mid-Atlantic Bight

Julia Levin[a], Hernan G. Arango[a], Bruce Laughlin[b], John Wilkin[a], Andrew M. Moore[b,*]

[a] *Department of Marine and Coastal Sciences, 71 Dudley Road, Rutgers University, New Brunswick NJ 08901, U.S.A..*
[b] *Department of Ocean Sciences, 1156 High Street, University of California, Santa Cruz CA 95062, U.S.A..*

Abstract

This paper explores the impact of the individual components of a coastal ocean observing system on estimates of the circulation derived from a state-of-the-art analysis and forecast system for the Mid-Atlantic Bight and Gulf of Maine. The foundation of these activities is the Regional Ocean Modeling System 4-dimensional variational (4D-Var) data assimilation platform, which is run in support of the Mid-Atlantic Regional Association Coastal Ocean Observing System as part of the U.S. Integrated Ocean Observing System. The specific focus of this study is on the impact of remote sensing observations from both space- and land-based platforms on estimates of cross-shelf transport in the vicinity of the National Science Foundation Ocean Observatories Initiative Pioneer array. Sea surface temperature (SST) and sea surface height (SSH) were found to have, on average, a similar impact on the transport estimates. However, during a typical 3-day 4D-Var assimilation cycle, approximately two orders of magnitude more observations of SST than SSH are used in the model, and closer analysis shows that each altimeter measurement has approximately 50 times more impact on the transport estimates than an individual SST observation. This highlights the value of altimetry data for ocean state estimation,

*Corresponding author
Email address: ammoore@ucsc.edu (Andrew M. Moore)

and the significance of expanding the altimeter constellation. The observations that are most impactful of all are *in situ* measurements of temperature and salinity, which have typically 3-4 times more impact than an individual SSH datum. A robust geographical distribution of the observation impacts emerges across a range of transport metrics which results from the combined influence of space-time dynamical interpolation and error covariance information within the 4D-Var system. The observation impact calculations suggest that High Frequency (HF) radar estimates of surface currents have relatively little direct influence on cross-shelf transport estimates. However, quantification of the sensitivity of these same estimates to changes in the observing system indicate that HF radar observations indirectly provide important information. This is understood in the current system by appealing to the idea of *borrowing strength* from the field of statistics in which some observations (satellite remote sensing in the case considered here) can borrow strength from other, seemingly less important observations.

Keywords: 4D-Var, observation impact, Mid Atlantic Bight

---

1. Introduction

Remote sensing observations from earth orbiting satellites are an essential and critical component of any ocean observing system, and are routinely assimilated in ocean models to constrain analyses and forecasts of the circulation. Even in a simple objective analysis system of the ocean state at a single time, quantifying the influence of any observing platform on the resulting circulation estimate can be challenging (e.g., Daley (1991) discusses various aspects of the interplay between observations in analyses). In a complex 4-Dimensional Variational (4D-Var) data assimilation system, unraveling the impact of individual observations collected through time can present a significant challenge because of the space-time dynamic interpolation embodied in the tangent linear and adjoint models, resulting in observation impacts that can be quite remote from the target region of interest.

The focus of this study will be on the impact of remotely sensed observations of sea surface temperature (SST), sea surface height (SSH) and surface currents in a near real-time analysis and forecast system of the Mid-Atlantic Bight (MAB) and the Gulf of Maine (see Fig. 1). The system is based on the Regional Ocean Modeling System (ROMS) and is run in support of the Mid-Atlantic Regional Association Coastal Ocean Observing System (MARACOOS) which forms part of the U.S. Integrated Ocean Observing System (IOOS). In addition to remote sensing observations, *in situ* observations from a variety of platforms are also assimilated into the model. However, the impact of *in situ* observations on the circulation estimates will be the subject of two companion studies.

The paper is organized as follows. The methodology used to compute the observation impacts is described in section 2, and the configurations of the Regional Ocean Modeling System (ROMS) and of the 4D-Var data assimilation system used here are summarized in sections 3 and 4 respectively. The method used to quantify the impact of observations on the 4D-Var analyses is predicated on the choice of specific indices or metrics that target particular aspects of the ocean circulation of interest. Section 5 introduces the metrics used in the present study, which are motivated by the ability of the observing system to shed light on cross-shelf transport processes. A quantitative assessment of the impact of the different components of the observing system on 4D-Var estimates of the cross-shelf transport is presented in section 6. Section 7 illustrates the role played by the various components of the data assimilation system in controlling the transfer of observational information throughout the estimation process. The synergy between observations from different platforms is demonstrated in section 8 by quantifying the sensitivity of the chosen metrics to changes in the observing system. A summary and some conclusions follows in section 9.

2. Observation Impacts and 4D-Var

The model used in this study is ROMS and the associated 4D-Var data assimilation platforms. The methodology used in ROMS to compute the impact of the observations on 4D-Var ocean circulation estimates is based on that employed routinely in numerical weather prediction and developed originally by Langland and Baker (2004, hereafter LB). The implementation in ROMS is described in detail by Moore et al. (2011c,a, 2017) so only a brief overview of the approach will be presented here.

Following the usual notation (Ide et al., 1997), the ROMS state-vector will be denoted by $\mathbf{x}$. If $\mathbf{x}^b$ denotes the background (or first-guess) then the analysis $\mathbf{x}^a$ resulting from data assimilation can be expressed as:

$$\mathbf{x}^a = \mathbf{x}^b + \mathbf{K}(\mathbf{y}^o - H(\mathbf{x}^b)) \tag{1}$$

where $\mathbf{y}^o$ denotes the vector of observations, H is the observation operator that maps from state-space to observation-space, and $\mathbf{K}$ is the Kálmán gain matrix. In the case of 4D-Var, the observation operator H includes the nonlinear model. The Kálmán gain matrix can be expressed in either the primal or dual form. Here we focus on the dual form since this is the flavor of 4D-Var that was used in the calculations described here. Thus, $\mathbf{K} = \mathbf{B}\mathbf{H}^T(\mathbf{H}\mathbf{B}\mathbf{H}^T + \mathbf{R})^{-1}$ where $\mathbf{B}$ and $\mathbf{R}$ are the background error and observation error covariance matrices respectively, and $\mathbf{H}$ represents the tangent linearization of the observation operator H. In 4D-Var $\mathbf{H}$ includes the tangent linearization of the nonlinear model and $\mathbf{H}^T$ includes the adjoint model.

Following the usual procedure (Talagrand and Courtier, 1987), the analysis $\mathbf{x}^a$ is identified by minimizing a cost-function expressed in the incremental form (Courtier et al., 1994). In ROMS, the cost function is minimized using the Lanczos formulation of the Restricted $\mathbf{B}$-Preconditioned Conjugate Gradient (RPCG) method (Gratton and Tshimanga, 2009) following Gürol et al. (2014). This is equivalent to a truncated Gauss-Newton method (Lawless et al., 2005) in which a sequence of linear minimization problems are solved itera-

tively (so-called inner-loops) where each sequence is linearized about an updated estimate of the circulation (so-called outer-loops). With this approach, the dual Kálmán gain matrix for each outer-loop can be factorized according to $\tilde{\mathbf{K}}_m = \mathbf{B}\mathbf{H}^T\mathbf{V}_m\mathbf{T}_m^{-1}\mathbf{V}_m^T\mathbf{H}\mathbf{B}\mathbf{H}^T\mathbf{R}^{-1}$, where m is the number of inner-loops and each of the m-columns of $\mathbf{V}_m$ represent the conjugate gradient descent directions normalized to unit amplitude (the so-called Lanczos vectors), and $\mathbf{T}_m$ is a known tridiagonal matrix. In this form, $\tilde{\mathbf{K}}_m$ represents a reduced-rank approximation of $\mathbf{K}$.

Following LB, the observation impacts are quantified in terms of their influence on a chosen scalar metric, $I(\mathbf{x})$. In particular, $\Delta I = I(\mathbf{x}^a) - I(\mathbf{x}^b)$ represents the change in I due to assimilating the observations $\mathbf{y}^o$. As shown by LB, to first-order $\Delta I \simeq (\mathbf{y}^o - H(\mathbf{x}^b))^T\mathbf{K}^T(\partial I/\partial\mathbf{x})|_{\mathbf{x}^b}$. The reduced-rank approximation $\tilde{\mathbf{K}}_m$ for $\mathbf{K}$ then leads to:

$$\Delta I \simeq (\mathbf{y}^o - H(\mathbf{x}^b))^T\mathbf{R}^{-1}\mathbf{H}\mathbf{B}\mathbf{H}^T\mathbf{V}_m\mathbf{T}_m^{-1}\mathbf{V}_m^T\mathbf{H}\mathbf{B}(\partial I/\partial\mathbf{x})|_{\mathbf{x}^b} \qquad (2)$$

where $(\partial I/\partial\mathbf{x})|_{\mathbf{x}^b}$ represents the derivative of I with respect to $\mathbf{x}$ evaluated using the background $\mathbf{x}^b$. Equation (2) shows that ΔI can be expressed as the dot-product of two vectors, the innovation vector $\mathbf{d} = (\mathbf{y}^o - H(\mathbf{x}^b))$ and the vector $\mathbf{g} = \mathbf{R}^{-1}\mathbf{H}\mathbf{B}\mathbf{H}^T\mathbf{V}_m\mathbf{T}_m^{-1}\mathbf{V}_m^T\mathbf{H}\mathbf{B}(\partial I/\partial\mathbf{x})|_{\mathbf{x}^b}$ which quantifies the impact of the observations on ΔI. Since each element of $\mathbf{d}$ is uniquely associated with an observation, so then are the corresponding elements of $\mathbf{g}$ such that the product $d_i g_i$ represents the contribution (aka *impact*) of the i^{th}-observation on ΔI. The observation impacts for a particular data assimilation cycle can, therefore, be easily computed from the archived 4D-Var Lanczos vectors (see also Trémolet, 2008).

3. ROMS Configuration

The configuration of ROMS used here is based on the MARACOOS grid shown in Fig. 1, and encompasses the MAB and the Gulf of Maine. The grid is primarily used to study the MAB circulation and cross-frontal exchange pro-

cesses between the continental shelf and the deep ocean. The model configuration is described in detail by Wilkin et al. (2018) and Levin et al. (2018) so only a brief description will be given here.

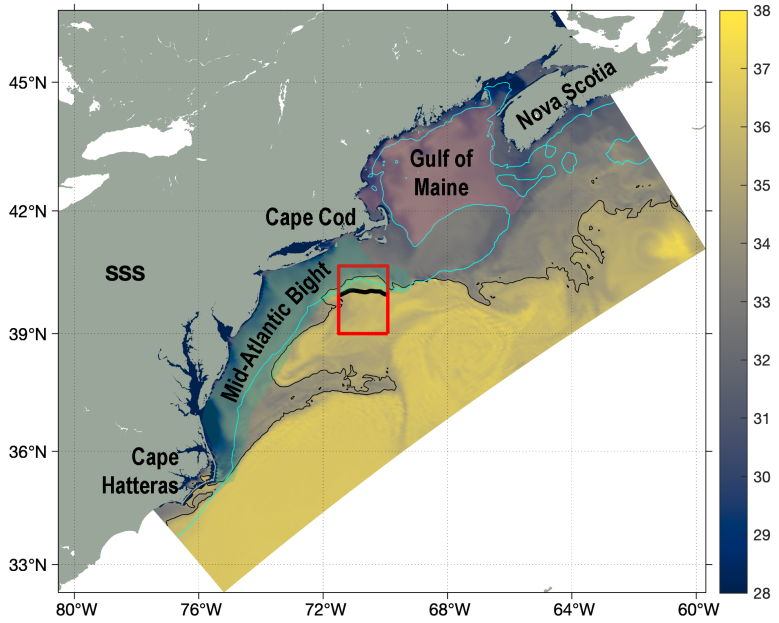

Figure 1: The ROMS model domain showing a 4D-Var sea surface salinity analysis on 16 May 2014. The black contour shows the 34.5 isohaline, and the thick black line denotes the location of target section of the 200 m isobath used to define the various transport metrics introduced in section 5. The red box indicates the position and the nominal extent of the OOI Pioneer glider observing array in relation to the target section. The cyan line is the 100 m isobath which can be used to delineate the Mid-Atlantic Bight and the Gulf of Maine as referred to in the main text.

The horizontal resolution is ~ 7 km and there are 40 vertical terrain-following levels stretched so that the thickness of the surface-most layer ranges from $0.1 - 1.8$ m and the bottom-most layer from $0.1 - 3.4$ m over the continental shelf to resolve boundary layer processes. Open boundary conditions use data from the Mercator-Océan global analysis (Drévillon et al., 2008) with temperature and salinity adjusted to remove seasonal bias compared to a local regional climatology Fleming (2016). The open boundary Mean Dynamic Topography

(MDT) and seasonal cycle of SSH variation were also adjusted for bias using a regional data assimilative climatological seasonal analysis computed following the procedure described by Levin et al. (2018) and Wilkin et al. (2018). The sub-tidal mesoscale variability captured by Mercator-Océan is retained. Harmonic tidal forcing (Mukai et al., 2002) was added to the boundary SSH and depth-averaged velocity data; in a domain this small the astronomical tide generating potential is negligible. Sea surface wind stress, and heat and freshwater fluxes were derived from 3-hourly National Centers for Environmental Prediction (NCEP) North American Mesoscale (NAM) forecast marine boundary layer conditions and standard bulk formulae Fairall et al. (2003). NAM air pressure is also imposed as a surface condition to the pressure gradient force so that the model computes a dynamic Inverted Barometer (IB) response. Accordingly, we add an equilibrium IB sea level term to the open boundary sea level data consistent with standard practice in altimeter data processing. Daily river inflows were imposed at 22 discharge sites based on U.S. Geological Survey and Water Survey of Canada observations and a statistical model that adjusts for ungauged portions of the watershed Wilkin et al. (2018).

4. ROMS 4D-Var

A full description of the 4D-Var configuration used here can be found in Levin et al. (2018) and Wilkin et al. (2018) so only a summary of the critical features will be presented here. The data assimilation approach employed is the dual formulation of the ROMS 4D-Var system (Moore et al., 2011c; Gürol et al., 2014) which was run using two outer-loops and seven inner-loops, a configuration that was adopted after extensive experimentation. Data were assimilated into the model using non-overlapping 3-day assimilation windows. The 4D-Var analysis $\mathbf{x}^a$ at the end of each 3-day assimilation window was used as the background $\mathbf{x}^b$ for the next assimilation cycle.

4.1. Observations

A summary of the data assimilated and the source of each data type that were available during the period Jan 2014 - Dec 2017 is given in Table 1. Figure 2a shows a time series of the total number of observations assimilated into the model during each 4D-Var cycle and indicates that $\sim 10^5$ observations are typically available during a 3-day window. There are generally fewer observations available during the winter due to a drop in the number of infra-red SST observations associated with increased cloudiness at this time of the year. A break down of the number of observations by type is shown in Fig. 2b and indicates that the lion's share of available data is in the form of satellite SST observations.

Type & platform	**Source**	**Sampling rate and resolution**	**Super-obs averaging**[1]	**Obs error**
AVHRR IR SST	MARACOOS.org & NOAA Coastwatch	4 passes per day, 1 km	3 hr	σ_b
GOES IR SST	NOAA Coastwatch	hourly, 6 km	3 hr	$2\sigma_b$
AMSR2, TRMM, and WindSat microwave SST	NASA JPL PODAAC	daily, 15 km		$1.25\sigma_b$
SSH: Jason, AltiKa, CryoSat	RADS, TU Delft	$\sim$ 1 pass daily, $\sim$ 4 km		0.04 m
in situ T, S: NDBC buoys, Argo floats, XBT, surface drifters	Met Office En4.2	variable[2]	standard levels[2]	$0.25\sigma_b\sigma_o/\sigma_{\max}$[3]
Surface currents: CODAR HF-radar	MARACOOS.org	hourly, 1 km	1 hr, 24 km	$0.5\sigma_b$
Glider T, S: MARACOOS	IOOS Glider DAC	variable[2]	standard levels[2]	$0.25\sigma_b\sigma/\sigma_{\max}$[3]
In situ T,S: GoM[5] buoys	NERACOOS.org[4]	hourly, 10 buoys	1 hr	σ_b
In situ u,v: GoM[5] buoys	NERACOOS.org[4]	hourly, 9 buoys[1]	1 hr	$0.5\sigma_b$

Table 1: A summary of the observational data assimilated into ROMS during 2014-2017, the procedure for forming super observations, and the observation errors assigned to each observation type. 1: All data that were sampled at a horizontal resolution higher than that of the model were formed into super observations at the resolution of the ROMS grid unless otherwise indicated. 2: Profile data were binned in the vertical using the WOD atlas standard depths (Boyer et al., 2009). 3: Here σ is the standard deviation of all observations that fall within a vertical bin (see comment 1) and $\sigma_{\max}$ is the maximum value of all σ in a vertical profile. 4: NERACOOS = North East Regional Association Coastal Ocean Observing System. 5: GoM=Gulf of Maine.

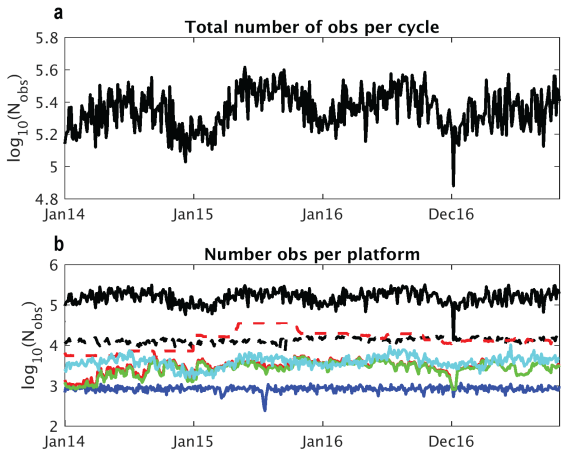

Figure 2: (a) A time series of $\log_{10}$ of the total number of observations assimilated into the model during each 3-day 4D-Var cycle. (b) Same as (a) but for the number observations of each type: SST - black; SSH - blue; *in situ* temperature - red; *in situ* salinity - green; *in situ* velocity - red dashed; gridded HF radar - black dashed; rejected obs - cyan.

It is standard procedure to thin the observations by combining multiple measurements of the same type that fall within a single grid cell, and that are closely spaced in time, into super observations. As noted in Table 1, all available SST observations were combined into 3-hourly fields and then, depending on the resolution of the dataset, either first formed into super observations or projected directly onto the model grid. Microwave and geostationary SST observations were used only in those grid cells that did not have polar orbiter infrared measurements. Altimetry observations were averaged if the tracks overlapped in space and time (within 7 km over a 2-hour interval).

Proximity analysis was performed on all observations to remove outliers. Quality control was also performed during each 4D-Var cycle following Järvinen and Undén (1997) as described by Moore et al. (2013). During this procedure, the innovation d_i associated with each observation is compared to the standard error based on the assumed standard deviations of the background (σ_b) and observation (σ_o) errors. Specifically, if $d_i^2 > \alpha^2(\sigma_b^2 + \sigma_o^2)$ then the observation is rejected and not included in the analysis. The threshold parameter α depends

on the observation type and is given in Table 2 for the analyses considered here. A time series of the total number of observations rejected during each 3-day 4D-Var cycle is shown in Fig. 2 and is typically $O(10^3)$ indicating that only $\sim 1\%$ of the total number of observations were rejected based on the criteria chosen.

The forward model computes the ocean‘s IB response and generates storm surge and coastal trapped waves within the domain, so we do not apply the Dynamic Atmosphere Correction (DAC) to altimeter SLA because that would introduce an inconsistency between observed and modeled SSH. The model also simulates the full tidal signal, but out of concern that small phase errors in the barotropic tide might dominate model-data sea level misfit we elect to de-tide the observations using the GOT4.10 correction of Ray (2013) and replace the tidal variability in the observations with a signal computed from a harmonic tidal analysis of a long free run of the MARACOOS model (see Wilkin et al., 2018). The same MDT described in section 3 applied to open boundary data is added to the augmented altimeter SLA to produce data for assimilation that has the best possible consistency between modeled and observed sea level signals.

Mesoscale sea level variability is correlated on time scales of a few to several days, but we presently lack an algorithm to formally impose this in the ROMS 4D-Var observation operator. Therefore, to encourage geostrophic adjustment of the dynamical fields and penalize the generation of barotropic surface gravity waves, we adopt the approach introduced by Zavala-Garay et al. (2014) whereby we assimilate each altimeter observation three times within a 2-hour time window: at the actual time of observation, and 1 hour before and 1 hour after. For these pseudo-observations the appropriate phase of the MARACOOS harmonic tide is added to the de-tided satellite value. We are unable to make a similar accommodation of the dynamic atmosphere effect so we inflate the observation error of the 1-hour lagged data to compensate.

4.2. Covariances

As described in Moore et al. (2011c), the 4D-Var background error covariance matrix $\mathbf{B}$ was modeled following the diffusion operator approach of Weaver and Courtier (2001). The decorrelation length scales assumed in $\mathbf{B}$ for errors in each state variable are listed in Table 2. Horizontal scales were estimated using a semi-variogram analysis (Banerjee et al., 2004) of the free-running model. In principle the horizontal and vertical decorrelation scales should be different as appropriate for each state variable, and also allowed to vary spatially and seasonally to reflect month-to-month changes in the circulation. Although larger vertical decorrelation scales help propagate information about the mesoscale from surface observations down through the water column, an overestimate of the vertical decorrelation scale in areas characterized by strong stratification can often drive spurious vertical oscillations. These oscillations are especially pronounced in the Gulf of Maine during summer months, eroding the stratification and pushing bottom temperature and salinity toward unrealistic values. Reliably estimating variations in covariance properties is very challenging (Moore et al., 2019), so as a compromise, we chose to use constant decorrelation scales based on the shortest vertical decorrelation scale, 10 m, estimated over the whole domain. While this may limit the influence of surface observations in the open ocean, the integrity of the summer time circulation in the MAB and GOM is preserved. Consequently, we are relying heavily on the tangent linear and adjoint model to propagate error information vertically via the contribution of $\mathbf{HBH}^T$ to the Kálmán gain. The standard deviations of the background errors were estimated from the mesoscale variability in an 11-year free run of the model.

The observation error covariance matrix $\mathbf{R}$ was assumed to be a diagonal matrix. The standard deviations for the observation errors were estimated by combining sensor error and representation error (from combining observations into super-observations). In some cases the observation errors were scaled by the background error standard deviation to make the model-data misfit reduction more uniform across different dynamic regimes. While these choices are

somewhat ad hoc, they were adopted to meet immediate demands to establish a real-time forecast instance of the 4D-Var system for MARACOOS. They work adequately in practice but deserve further evaluation and refinement. The observation errors assigned to the various platforms, and any associated rescalings, are summarized in Table 1.

State variable	**Horizontal decorrelation scale (km)**	**Background quality control parameter α**
SSH	40	5
Velocity	40	3
Temperature	40	6
Salinity	15	6

Table 2: A summary of the decorrelation scales assumed for background errors in each state variable. The vertical decorrelation length scale for all state variables was chosen to be 10 m. The parameter α used for the background quality control rejection criteria is also indicated.

4.3. 4D-Var diagnostics

The performance of the 4D-Var system is described by Levin et al. (2018) and Wilkin et al. (2018), so only a summary is presented here.

Figure 3 shows probability density functions (pdfs) for the innovations $\mathbf{d} = (\mathbf{y}^o - H(\mathbf{x}^b))$ associated with observations of SST, SSH, *in situ* temperature and *in situ* salinity from the 1st outer-loop of all 4D-Var cycles. If $\mathbf{B}$ and $\mathbf{R}$ are correctly prescribed, the innovations $\mathbf{d}$ should be normally distributed with a covariance given by $(\mathbf{HBH}^T + \mathbf{R})$. Therefore, for reference, Fig. 3 also shows the pdfs for normal distributions with the same mean and standard deviation as the innovations computed during the 1st outer-loop. A comparison of the normal distribution curves with the innovation pdfs in Fig. 3 indicates that, for all observation types, the innovation pdfs depart significantly from the expected normal distributions, and are more reminiscent of Laplace distributions. For the most part, the mean innovations for SST, *in situ* temperature and *in situ* salinity

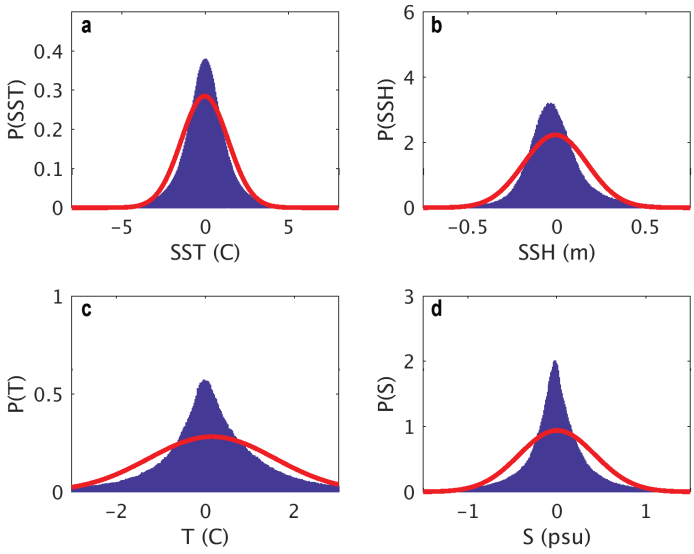

Figure 3: Probability density functions (pdfs) of the 4D-Var increments in (a) SST, (b) SSH, (c) *in situ* temperature, and (d) *in situ* salinity computed from all 4D-Var cycles spanning the 2014-2017 period.

are close to zero. The mean SSH innovation, however, is negative indicating that, on average, the mean model SSH exceeds that observed. During the 2nd outer-loop, the innovation pdfs are qualitatively similar to those for the 1st outer-loop (not shown).

A very simple measure of the performance of the 4D-Var system can be obtained by inspecting the contribution of the observations $J_o = (\mathbf{y}^o - \mathbf{y}^m)^T \mathbf{R}^{-1} (\mathbf{y}^o - \mathbf{y}^m)$ to total the 4D-Var cost function, where $\mathbf{y}^m$ denotes the model circulation estimate evaluated at the observation points. Figure 4 shows the time series of the ratio of the final to initial values of J_o at the end of each outer-loop associated with SSH and temperature observations. During both outer-loops, the average fit of the model to the observations is improved by the 4D-Var procedure.

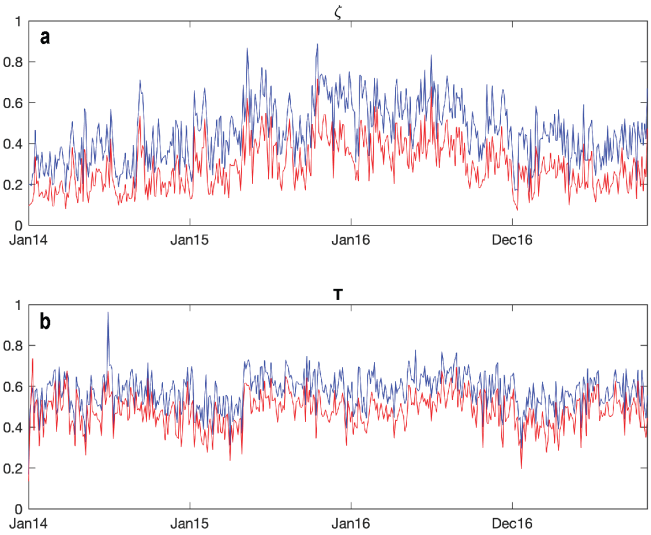

Figure 4: Time series of the ratio of final to initial values of J_o for the 1st outer-loop (blue) and the 2nd outer-loop (red) for the contributions from observations of (a) SSH and (b) SST.

5. Circulation Metrics

An important circulation feature of the northeast shelf is the MAB shelf-break front (Linder and Gawarkiewicz, 1998), which separates the warm, saline water of the Gulf Stream from the cooler, fresher waters of the continental slope. The exchange of water masses across the continental shelf-break is of considerable interest and is influenced by a range of different processes (Gawarkiewicz et al., 2018). One such process is illustrated in Fig. 1 and is associated with the intrusion of Gulf Stream rings onto the continental shelf as captured by the 4D-Var analysis on 16 May 2014. This event was studied in detail by Zhang and Gawarkiewicz (2015) and takes the form of a streamer of high salinity water associated with a warm core eddy impinging on the shelf. A series of metrics, I, are considered here to quantify the impact of the observations on the ability of the model 4D-Var analyses to capture these kinds of events. Specifically, the metrics chosen evaluate the magnitude of the cross-shelf exchange fluxes in the

vicinity of the OOI Pioneer Array (cf Fig. 1). This region was chosen because it has been the subject of other studies reported in the literature (e.g. Garvine et al., 1989; Linder and Gawarkiewicz, 1998; Chen and He, 2010), and also because we wish to quantify the relative impact of remotely sensed observations and *in situ* observations from the densely sampled Pioneer array. The impact of the *in situ* observing systems will be explored in a companion study and not reported on in detail here.

The following metrics were used to quantify the impact of the observations on the 4D-Var estimates of cross-shelf exchange at the shelf-break in the vicinity of the OOI Pioneer array:

$$I_u = \int_s \int_{-h}^{0} (\overline{u}_n - \tilde{u}_n) dz ds \tag{3}$$

$$I_{uT} = (\rho_0 c_p A)^{-1} \int_s \int_{-h}^{0} (\overline{u}_n - \tilde{u}_n)(\overline{T} - \tilde{T}) dz ds \tag{4}$$

$$I_{uS} = (10^3 \rho_0 A)^{-1} \int_s \int_{-h}^{0} (\overline{u}_n - \tilde{u}_n)(\overline{S} - \tilde{S}) dz ds. \tag{5}$$

In each case, $\int_s \ldots ds$ represents an integral along a vertical section following a segment of the $h = 200$ m isobath, which is nominally identified as the location of the continental shelf-break, and A is the total cross-sectional area. The location of this vertical section is indicated in Fig. 1 and cuts through the middle of the region sampled by the Pioneer glider array. In (3) - (5), u_n corresponds to the component of the velocity that is locally normal to the section; $\tilde{u}_n$, $\tilde{T}$ and $\tilde{S}$ represent the mean seasonal cycle; and an over-bar denotes the time average over the duration of the data assimilation window. Therefore, I_u, I_{uT} and I_{uS} represent measures of the 4D-Var cycle average total volume transport, heat transport and salt transport respectively crossing the shelf-break.

Figures 5a-c show time series of each metric computed from the background circulation estimates. Also shown are time series of each metric calculated from a free run of the model without data assimilation. The cross-shelf volume transport I_u (Fig. 5a) varies in the range $\sim \pm 7$ Sv during the 4D-Var analyses and $\sim \pm 5$ Sv during the free run. Some features of the transport variability are similar in both cases, suggesting that they are associated mainly with surface

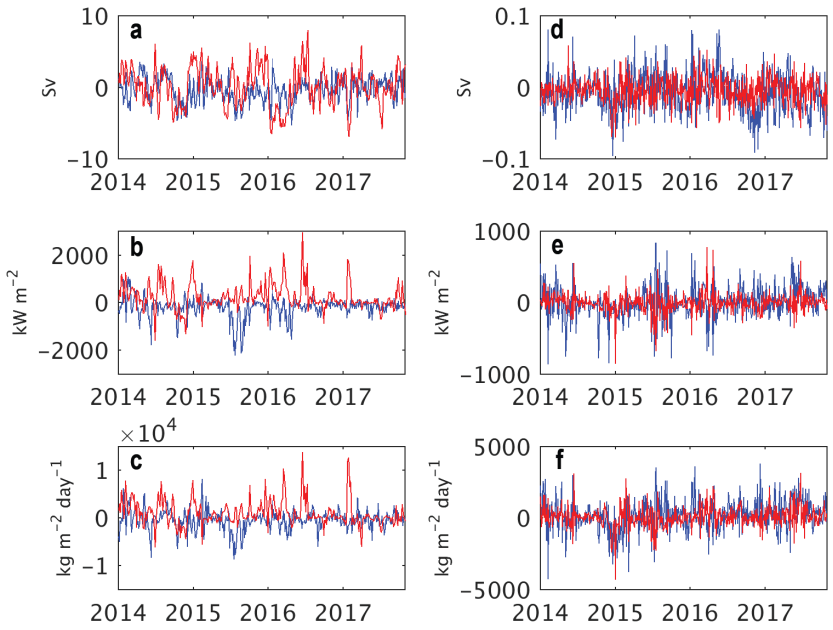

Figure 5: Time series of the circulation metrics computed from the 4D-Var analyses (red) and from a free running without data assimilation (blue) for (a) I_u, (b) I_{uT} and (c) I_{uS}. Also shown are time series of the circulation metric increments at the end of the 1st outer-loop (blue) and the 2nd outer-loop (red) for (d) ΔI_u, (e) ΔI_{uT} and (f) ΔI_{uS}.

forcing, although in general the two-time series are quite different ($r = 0.3$). The variability in cross-shelf heat transport I_{uT} (Fig. 5b) and cross-shelf salt transport I_{uS} (Fig. 5c) is very different in the 4D-Var analyses and the free run. Even though I_u is generally in the same direction in both cases, I_{uT} is often in the opposite direction in the analysis and free run, which is due to the temperature anomalies in the two cases being of opposite sign (not shown). Similarly for the time series of I_{uS}. Clearly then, data assimilation yields significant corrections to the cross-shelf exchange of water masses.

The impact of the observations on each metric was quantified according to (2) which represents a first-order linearization of the increment in each metric arising from data assimilation during each cycle. Equation (2) shows that an im-

portant ingredient of these calculations is $(\partial I/\partial \mathbf{x})|_{\mathbf{x}^b}$. Since (4) and (5) present nonlinear metrics, computation of this first derivative represents an additional linear approximation in the procedure.

Since two outer-loops are employed in the 4D-Var analyses, it is necessary to compute the observation impacts separately for each outer-loop. If $\mathbf{x}_1^a$ denotes the 4D-Var analysis at the end of the 1st outer-loop, then the observation impacts would be quantified according to $\Delta I_1 \simeq (\mathbf{y}^o - H(\mathbf{x}^b))^T \tilde{\mathbf{K}}_1^T (\partial I/\partial \mathbf{x})|_{\mathbf{x}^b}$ and $\Delta I_2 \simeq (\mathbf{y}^o - H(\mathbf{x}_1^a))^T \tilde{\mathbf{K}}_2^T (\partial I/\partial \mathbf{x})|_{\mathbf{x}_1^a}$. Here, ΔI_1 and ΔI_2 represent the increments in the metric I at the end of the 1st and 2nd outer-loop respectively, $\tilde{\mathbf{K}}_1$ and $\tilde{\mathbf{K}}_2$ are the reduced-rank Kálmán gain matrices for each outer-loop, and $(\partial I/\partial \mathbf{x})|_{\mathbf{x}^b}$ and $(\partial I/\partial \mathbf{x})|_{\mathbf{x}_1^a}$ represent the derivatives of the metric I evaluated using $\mathbf{x}^b$ and $\mathbf{x}_1^a$. However, $\mathbf{x}_1^a$ and $\tilde{\mathbf{K}}_2$ are no longer independent of the observations, in which case the expression for ΔI_2 does not provide an unambiguous indicator of the impact of each datum during the 2nd outer-loop. Trémolet (2008) discusses this issue in detail and notes that the observations typically exert the most influence on the analysis during the 1st outer-loop, and as such, their impacts during the 1st outer-loop are useful as an overall indicator since almost always $\Delta I_1 > \Delta I_2$, which is the case here. This is illustrated in Figs. 5d-f which shows time series of ΔI_1 and ΔI_2 associated with each metric computed using the tangent linear approximation (2). The increments in each index computed directly from the difference between the non-linear model trajectories $\mathbf{x}_1^a$ and $\mathbf{x}^b$ are almost indistinguishable from those using the tangent linear assumption as evidenced by the high correlation coefficients (0.98 for ΔI_u and 0.91 for ΔI_{uT} and ΔI_{uS}). Therefore, the tangent linear approximation, on which the following analysis of the observation impacts is based, is reliable. Figure 5d shows that 4D-Var makes relatively modest increments in cross-shelf volume transport ΔI_u compared to the total value of Fig. 5a, a desirable feature, since this indicates that the mass transport is not undergoing significant and/or abrupt changes as a result of assimilating the observations. The increments in cross-shelf heat and salt transport (Figs. 5e and 5f), however, are a much more significant fraction of the total transports (Figs. 5b and 5c) indicative of correc-

tions in the temperature and salinity through the water column as informed by the observations.

6. Observation Impacts

Using (2), the increments ΔI_u, ΔI_{uT}, and ΔI_{uS} were decomposed into the contributions of each component of the observing system. Figure 6 shows a summary of the root mean square (RMS) impact, averaged over all assimilation cycles for each circulation metric, of the various and different types of observing platforms. While *in situ* temperature and salinity observations clearly exert the most influence on each circulation metric, altimetry and SST also have significant impact. Figure 6 also shows the RMS impact of each altimeter and SST platform evaluated separately, and the repeated assimilation of SSH observations described in section 4.1 has been accounted for here and in all subsequent figures. The RMS impacts across the different altimeter platforms are similar, while for SST, it is observations from AVHRR and AMSR that have the most influence. This possibly reflects that they are quite complementary observations, with infrared AVHRR having high resolution but incomplete spatial coverage due to clouds, whereas microwave AMSR delivers data in all weather but at reduced resolution and accuracy. The impacts for these platforms also scale inversely with the observation errors that are assigned to each. In the case of altimetry, the observation errors assumed for each platform are the same, while for infrared SST, AVHRR errors are half those of GOES (see Table 1). Remotely sensed estimates of surface current derived from shore-based HF radars have only modest impact on the three outer shelf transport metrics we have chosen to consider here.

6.1. Impacts per datum

The RMS impact per datum on each transport metric of each observation type averaged over all 4D-Var cycles is shown in Table 3. When viewed in this way, the least impactful observations from a *per datum* perspective are SST observations. However, on aggregate, because there are so many SST observations

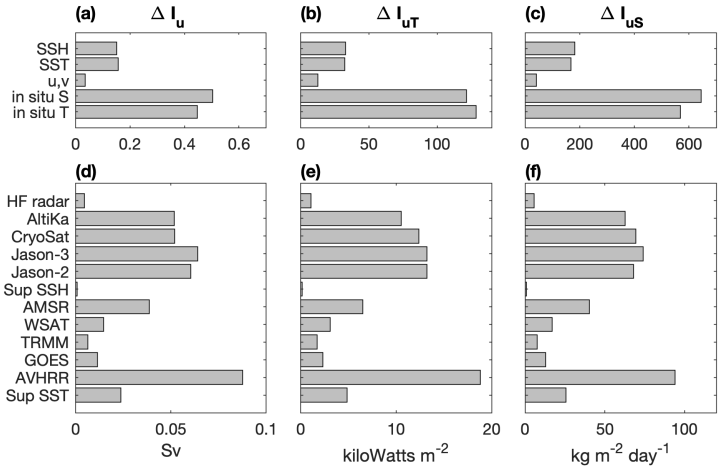

Figure 6: Histograms of the RMS impact averaged over all 4D-Var cycles for observations of different types (a-c) and individually for different remote sensing platforms (d-f) for the increments in each of the metrics ΔI_u, ΔI_{uT} and ΔI_{uS}. The repeated assimilation of SSH observations described in section 4.1 has been accounted for here and in all subsequent figures. Sup SSH and Sup SST refer to super-observations for SSH and SST respectively.

(cf Fig. 2), SST observations contribute significantly to the cross-shelf transport estimates (Fig. 6). Table 3 shows that a single altimeter observation is ~ 50 times more impactful than a single SST observation, and this is generally true across all altimetry platforms. The data that have the greatest impact are *in situ* temperature and salinity observations which are $\sim 3-4$ times more impactful than individual altimeter measurements. Individual *in situ* measurements of velocity and HF radar surface current estimates have a similar impact to SST.

6.2. Impact as an indicator of 4D-Var performance

As noted in section 2, 4D-Var is an iterative algorithm, and it is of interest to quantify the relative impact of each observation type during the iteration procedure since this provides a useful indicator of 4D-Var performance. To this end, Fig. 7 shows the RMS impact on each transport metric of the different observation types computed during each inner-loop for both outer-loops and averaged over all 4D-Var cycles. The dominant impact of *in situ* temperature

	I_u (m^3 s^{-1})	I_{uT} (W m^{-2})	I_{uS} (10^{-3} kg m^{-2} day^{-1})
SST	5	1.2	6
SSH (all)	250	54	310
Jason-2	240	54	280
Jason-3	320	63	380
AltiKa	240	50	310
Cryosat	240	54	300
in situ T	780	200	930
in situ S	1000	262	1230
in situ u,v	4	1.2	5
HF radar	2	0.6	3

Table 3: The RMS impact per datum of each observation type and altimeter platform on the three transport metrics averaged over all 4D-Var cycles. The repeated assimilation of SSH observations described in section 4.1 has been accounted for here.

and salinity observations is again very evident in each case, in keeping with Fig. 6. Figure 7 shows, however, that on average SST and SSH observations collectively have an almost equal average impact. As noted in section 5, the average impact of a given observation type is typically larger during the 1st outer-loop than during the 2nd outer-loop, in agreement with the arguments put forth by Trémolet (2008).

An interesting feature of Fig. 7 is the rate at which the RMS impacts of each observation type asymptote to a near constant value. In the case of *in situ* hydrographic observations, the RMS impacts reach their near maximum value after just three inner-loops, indicating that much of the useful information has been extracted from these data early on in the assimilation procedure. In contrast, the RMS impacts for SST and SSH continue to exhibit an upward trend at the end of each outer-loop, suggesting that there is further useful information that could be extracted from these observations by 4D-Var. While it is by no means obvious how the 4D-Var system should be tuned or reconfig-

ured to capitalize on this additional information, it is a topic worthy of further investigation.

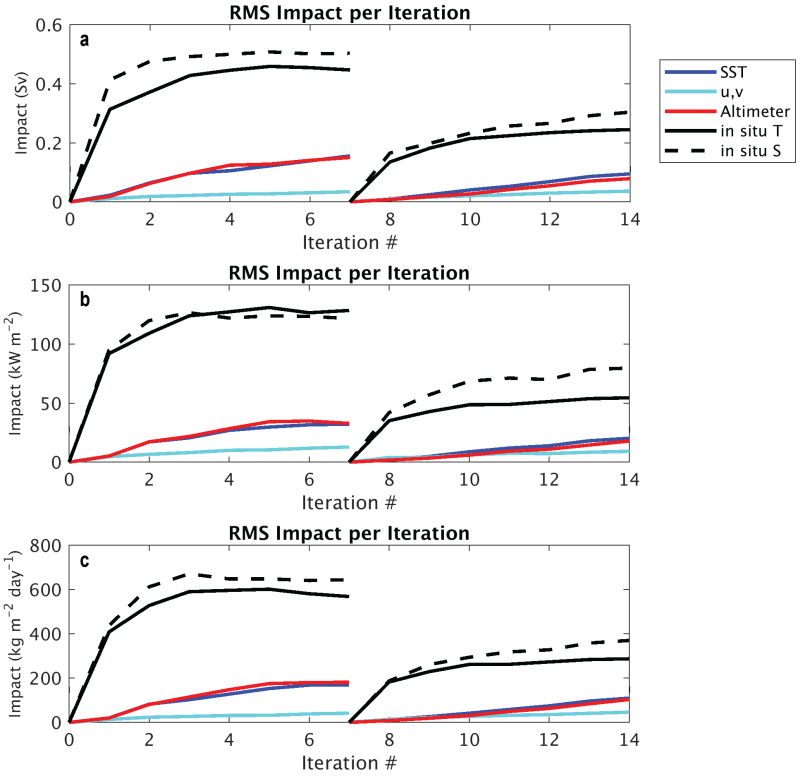

Figure 7: The RMS impact of each data type during each inner-loop averaged over all assimilation cycles for (a) ΔI_u, (b) ΔI_{uT} and (c) ΔI_{uS}. Iterations 1-7 correspond to the 1st outer-loop, while iterations 8-14 correspond to the 2nd outer-loop.

6.3. Impact and bias

Another illuminating summary of the observation impacts for different observation types is in the form of scatter plots and 2-dimensional histograms of the impact per observation versus the innovation. Examples are shown in Fig. 8 for the cross-shelf volume transport increments ΔI_u for the 1st outer-loop from all assimilation cycles. Grey dots are the innovation-impact pairs for every observation, while the colors summarize the frequency of these pairs in 2-dimensional innovation-impact bins. Red dashed lines divide each plot into

a contingency diagram and the number of points in each quadrant is noted in italics.

Focusing of the colored 2-D histogram data, the plots in Fig. 8 generally take the form of "butterfly wings" with small innovations (for which the difference between the background and the observation is small) having little impact on ΔI_u, but as the magnitude of the innovations increase the size of the impacts also increase; reassuringly however, observations associated with very large innovations (large model-data differences) generally exert lesser impact on ΔI_u, which is evident from the tapering of the "wings". This is a desirable feature of the data assimilation system because these are quite likely cases of observations that passed the quality control threshold through the coincidence of a poor background solution at a bad observation location, and we would not want these to adversely impact the analysis.

It is altimeter observations that have the greatest tendency to depart from the butterfly-wing pattern with streaks of dots emanating from the plot origin indicating high impact for modest innovation. However, these do not register in the 2-D histograms indicating that they stem from a handful of points only. We can report that upon deeper inspection we find that each streak is typically associated with a single satellite pass, indicating either altimeter data with correlated errors along-track that pass quality control or, alternatively, short-lived sea-level variability events that were wholly missed by the model yet captured by satellite. We have not included similar plots for individual SST platforms, but they are all qualitatively similar to Fig. 8b.

In the absence of bias, the four quadrants of a contingency diagram should be approximately equally populated. This is generally the case for temperature, salinity and velocity (Figs. 8b-e) but for SSH (Fig. 8a) there is a moderate preponderance of negative innovations, which is consistent with the pdf in Fig. 3b. This bias occurs across all altimeters (Figs. 8f-i) which could indicate that either or both of the Mean Sea Surface (MSS) to which SLA is referenced, or the Mean Dynamic Topography (MDT) that is added to SLA, tend to be too high. In shelf waters the MSS_CNES_CLS_15 product (Schaeffer et al., 2016)

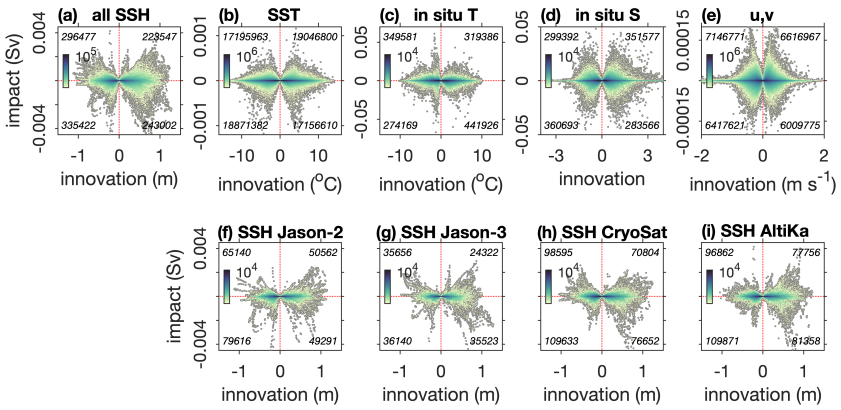

Figure 8: Scatter plots (dots) with superimposed 2-D histograms (color; log scale) of impact on ΔI_u versus innovation for different observation types (a-e), and for each individual altimeter satellite (f-i). Red dashed lines divide each plot into four quadrants of a contingency diagram, with the number of points that populate each quadrant indicated.

is an extrapolation by optimal interpolation of multi-mission satellite altimeter data derived using standard range and geodetic corrections that exclude shallow water observations. This introduces errors that we cannot easily quantify but would not expect to be consistently biased. While our MDT is a significant improvement over the AVISO standard product (Levin et al., 2018), if the mean along-shelf current is still less than in reality this would be consistent with a negative bias in innovations.

The results for each altimeter satellite (Figs. 8f-i) indicate that positive innovations typically have the largest impact on ΔI_u. We presently do not have a hypothesis for what characteristic of the model, the data, or their respective errors, might account for this. Positive innovations also have moderately larger

impact for ΔI_{uT} and ΔI_{uS} metrics (not shown). Data assimilation has thus highlighted potential issues with altimeter data bias deserving further analysis.

6.4. Spatial variations in impact

The geographic variations in the impact associated with each observation type are a reflection of the dynamical processes that control the flow of information through the data assimilation system. With this in mind, Fig. 9 shows the geographic variability of the RMS impact (for all three transport metrics) of SST and SSH observations that fall within each surface grid cell of the model for the 1st outer-loop for all 4D-Var cycles. It is important to note that the RMS impacts in Fig. 9 are plotted on a log-scale, and indicate that the average impact on each metric within the Pioneer target region of observations from different parts of the model domain can vary by 2-3 orders of magnitude. Figure 9 exhibits several remarkable features. First, it is evident that observations both close to and remote from the target area can have a significant impact on each transport metric in the case of both SST and SSH. Second, for a given observation type, the geographic distribution of the regions of high and low impacts display remarkable qualitative similarities across all three metrics. In the case of SST (Figs. 9d-f), observations in the vicinity of the target section have a significant impact as one might expect. However, observations further afield, such as near the north wall of the Gulf Stream, have an equally significant impact. The impact of observations that are remote from the target region are an indication of the critical role played by ocean dynamics in propagating information through the 4D-Var system. In the case of altimetry (Figs. 9a-c), SSH observations in the vicinity of the target section have a relatively small impact on each metric, and instead, it is observations that are remote from the target area that exert the largest influence. In particular, observations from altimeter passes over the Gulf Stream and some near-shore regions such as western Long Island and the Gulf of Maine have the largest impact. In the case of ΔI_u, these are the data that populate the wings of the scatter plots in Figs. 8f-i.

The RMS impacts associated with each altimeter and SST platform are

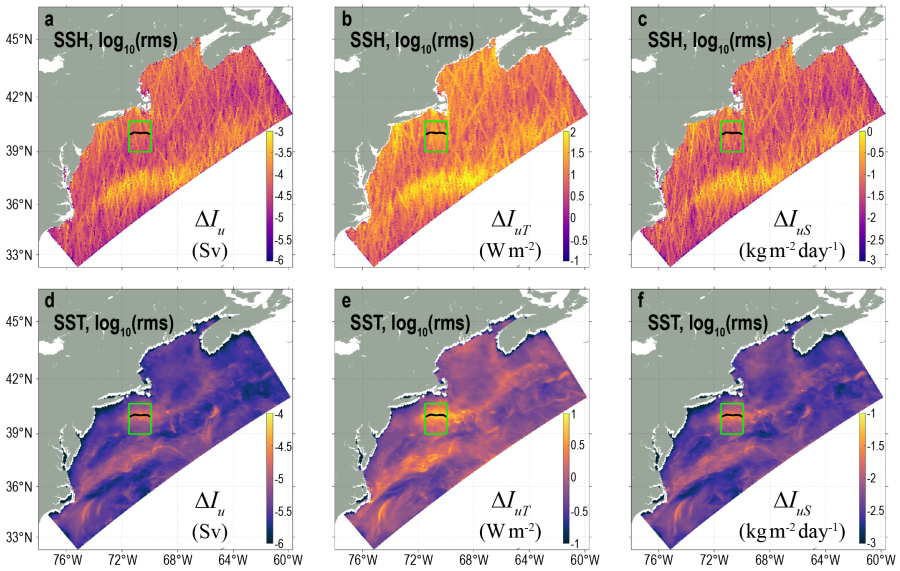

Figure 9: The spatial distribution of $\log_{10}$ of the RMS impact of observations falling within each model grid cell averaged over the all 4D-Var cycles (2014-2017) for altimeter observations and for (a) ΔI_u , (b) ΔI_{uT} and (c) ΔI_{uS}. The RMS impact of satellite SST observations is also shown for (d) ΔI_u , (e) ΔI_{uT} and (f) ΔI_{uS}. In each case the RMS is shown at the end of the 1st outer-loop. The location of the 200 m isobath target section crossing the Pioneer glider array (green box) is indicated.

similar to those of Fig. 9, and the patterns are robust from year-to-year (not shown). The origin of these robust patterns will be explored next.

7. The Mechanics of Observation Impacts

7.1. An Illustrative Example

Before examining the nature of the geographic distributions of observation impact discussed in section 6.4, it is instructive to pause and consider the underlying mechanics of the calculations described by (2), since this can guide our dynamical interpretation of the results. With this in mind, the circulation metrics introduced in section 5 can be written generically in a discrete form as either $I_L = \sum_{i=i_s}^{i_e} \mathbf{h}_i^T \mathbf{x}_i$ or $I_Q = \sum_{i=i_s}^{i_e} \mathbf{x}_i^T \mathbf{E}_i^T \mathbf{x}_i$ where $\mathbf{x}_i = \mathbf{x}(t) = \mathbf{x}(i\Delta t)$ is the state-vector at time $t = i\Delta t$ with time step Δt, and i_s and i_e denote the

start and end time steps of the integration, the assimilation window in the cases considered in section 5. For the class of linear metrics I_L and quadratic metrics I_Q, the elements of the vector $\mathbf{h}_i$ and matrix $\mathbf{E}_i$ represent weights appropriate for the metric. Specifically, I_u belongs to the class I_L while I_{uT} and I_{uS} belong to the class I_Q. Regarding (2), the derivative of each metric with respect to $\mathbf{x}$ can be expressed as:

$$\partial I_L/\partial \mathbf{x} \simeq \sum_{i=i_s}^{i_e} \mathbf{M}_i^T \mathbf{h}_i \tag{6}$$

$$\partial I_Q/\partial \mathbf{x} \simeq \sum_{i=i_s}^{i_e} \mathbf{M}_i^T (\mathbf{E}_i + \mathbf{E}_i^T)\mathbf{x}_i \tag{7}$$

where $\mathbf{M}_i^T$ is the adjoint model linearized about $\mathbf{x}_i$, in which case the increment in each metric during the 1st outer-loop can be expressed as:

$$\Delta I_L \simeq (\mathbf{y}^o - H(\mathbf{x}^b))^T (\mathbf{HBH}^T + \mathbf{R})^{-1} \mathbf{HB} \sum_{i=i_s}^{i_e} \mathbf{M}_i^T \mathbf{h}_i \tag{8}$$

$$\Delta I_Q \simeq (\mathbf{y}^o - H(\mathbf{x}^b))^T (\mathbf{HBH}^T + \mathbf{R})^{-1} \mathbf{HB} \sum_{i=i_s}^{i_e} \mathbf{M}_i^T (\mathbf{E}_i + \mathbf{E}_i^T)\mathbf{x}_i^b. \tag{9}$$

Similar expressions hold for outer-loop n if $\mathbf{x}^b$ is replaced by $\mathbf{x}_{n-1}^a$.

Figure 10 shows a schematic illustration of the various matrix operations in (8) and (9) for the case of an idealized steady zonal shear flow. The example illustrated is for a metric I defined at a single instant in time along a target section that is perpendicular to the jet axis (Fig. 10a). In this case, the derivative $\partial I/\partial \mathbf{x}|_{\mathbf{x}^b}$ represents a line of impulses at each grid point that contributes to I. When multiplied by the adjoint propagator $\mathbf{M}^T$, Fig. 10b illustrates how each impulse is integrated backwards in time and is thus advected upstream against the shear flow. Part of the adjustment process will involve the generation of waves which, in the adjoint model, will move in the opposite direction to their counterparts in the tangent linear model. At this point, the adjoint model solution represents a weighted sum of the Green‘s functions associated with all the grid points that contribute to I. The next operation in the sequence (Fig. 10c) is a multiplication by the background error covariance matrix $\mathbf{B}$ which acts to

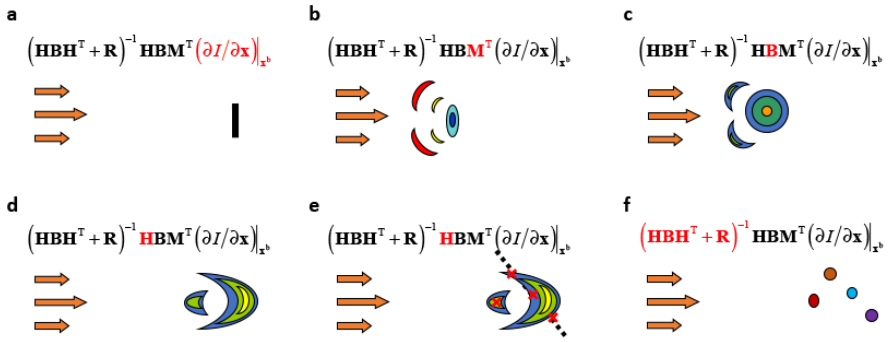

Figure 10: A schematic representation of the outcome of each operation represented by equations (8) or (9) for the case of an idealized shear flow which is indicated by the arrows. The metric I used is defined by all the grid points that lie along the black line in (a) that is orthogonal to the jet axis. The operation considered in each case is highlighted in red in the accompanying equation. In (e), the ground track of an idealized satellite pass (black dashed line) and some associated observations (red crosses) is shown, as well as the site of an isolated *in situ* observation.

spread information in space in a way that is consistent with the correlation lengths associated with the prior uncertainties. As illustrated in Fig. 10c, this operation has the effect of smoothing the Green's functions. It is followed by the operation of $\mathbf{H}$; the tangent linear model sampled at the observations. In the example considered here, the smoothed Green's functions are subsequently advected downstream and deformed by the shear flow (Fig. 10d). The resulting tangent linear fields are then sampled at the observation points (Fig. 10e) where, for simplicity, we show the ground track of an idealized satellite pass and an isolated *in situ* observation. Finally, the tangent linear fields at the observation points are multiplied by the inverse stabilized representer matrix $(\mathbf{HBH}^T + \mathbf{R})^{-1}$ (Fig. 10f), an operation that effectively removes the total prior covariances between the observation points. The product of each resulting element of the solution in Fig. 10f with the corresponding innovation will yield the contribution of each observation in Fig. 10e to ΔI (*i.e.*, the impact on I).

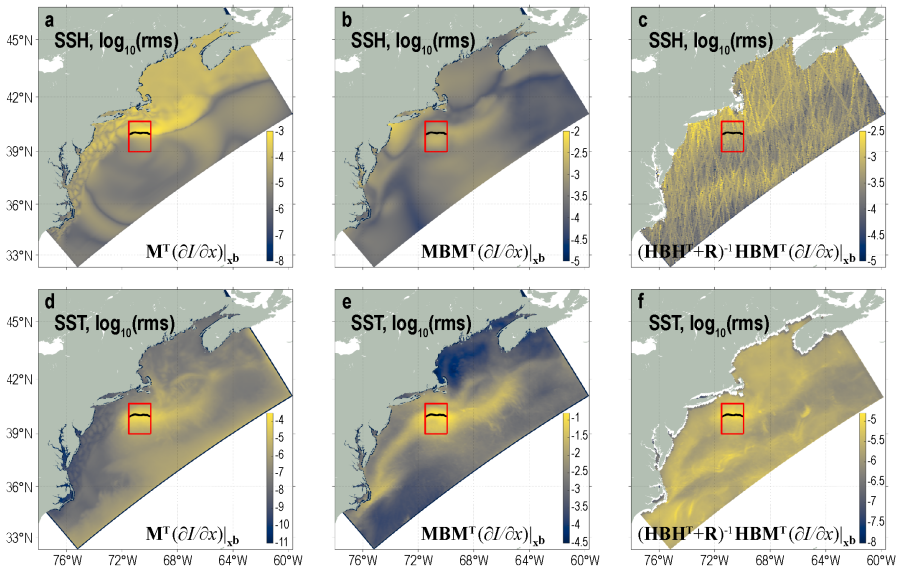

Figure 11: The spatial distributions of $\log_{10}$ RMS for SSH (a-c) and SST (d-f) resulting from the operations indicated by the accompanying equations, averaged over all 4D-Var cycles (2014-2017). The target section of the 200 m isobath for each metric is shown for reference crossing the Pioneer glider array (red box).

7.2. MAB Analyses

The different steps of the observation impact calculations illustrated schematically in Fig. 10 as applied to the MAB cross-shelf transport metric I_u are shown in Fig. 11 to elucidate which steps in the procedure are predominantly responsible for the robust patterns of SST and SSH impact that are apparent in Fig. 9. Specifically, the RMS result of the different stages of the observation impact procedure associated with observations of SSH and SST are shown in Fig. 11 averaged over all 4D-Var cycles spanning the full period 2014-2017. Fig. 11a shows the RMS of the derivative (6) for I_u with respect to SSH, the weighted sum of Green's functions associated with all of the model grid points (in space and time) that contribute to I_u (analogous to Fig. 10b). The largest values are located in the immediate vicinity of the target section, although there are significant upstream influences associated with advection by the equatorward shelf-break jet. There are also localized high values inshore of the target section

that are most likely associated with wave propagation. Although significantly lower in amplitude, the RMS of the Green's functions also reveal deep water impacts of SSH observations. The RMS SST component of (6) for I_u is illustrated in Fig. 11d and, like SSH, exhibits both local and remote influences of SST on the target section volume transport. In addition to the upstream influences of the shelf-break jet, weak influences of SST in the vicinity of the open boundaries are also evident.

The influence of smoothing by $\mathbf{B}$ and the propagation of the smoothed Green's functions forward in time by the tangent linear model (analogous to Fig. 10d) is illustrated in Figs. 11b and 11e for the SSH and SST components. To illustrate first the propagation of information in the full state-space of the model, Figs. 11b and 11e show the result of operating with $\mathbf{M}$ rather than $\mathbf{H}$ (i.e the tangent linear model solution before it is sampled at the observation locations). For SSH, Fig. 11b shows that much of the signal present in Fig. 11a is advected or propagated away from the target region, or dissipated, with some of the largest values along the New Jersey coast. In the case of the SST component, Fig 11e indicates that much of the impact remains close to the target section, although upstream and downstream influences of the shelf-break jet are apparent also.

Figures 11c and 11f show the result of sampling the tangent linear model solution at the observation points and removing the total prior covariance between the observation locations (cf Figs. 10e and10f). Only at this stage of the observation impact calculations do we begin to see the emergence of the robust geographic patterns in the RMS impact fields that are so evident in Fig. 9. This indicates that some or all of this structure is associated with the inverse stabilized representer matrix. Recall though that according to (8) the impact of each observation on I_u is given by the product of the innovation of each observation with the corresponding elements of Figs. 11c and 11f in the case of SSH and SST. With this in mind, Fig. 12 shows the RMS of the SST and SSH elements of the innovation vector $\mathbf{d}$, and the RMS of $(\mathbf{HBH}^T + \mathbf{R})^{-1}\mathbf{d}$ associated with the SST and SSH observations. A comparison of Figs. 12a and 12b with Figs. 9a

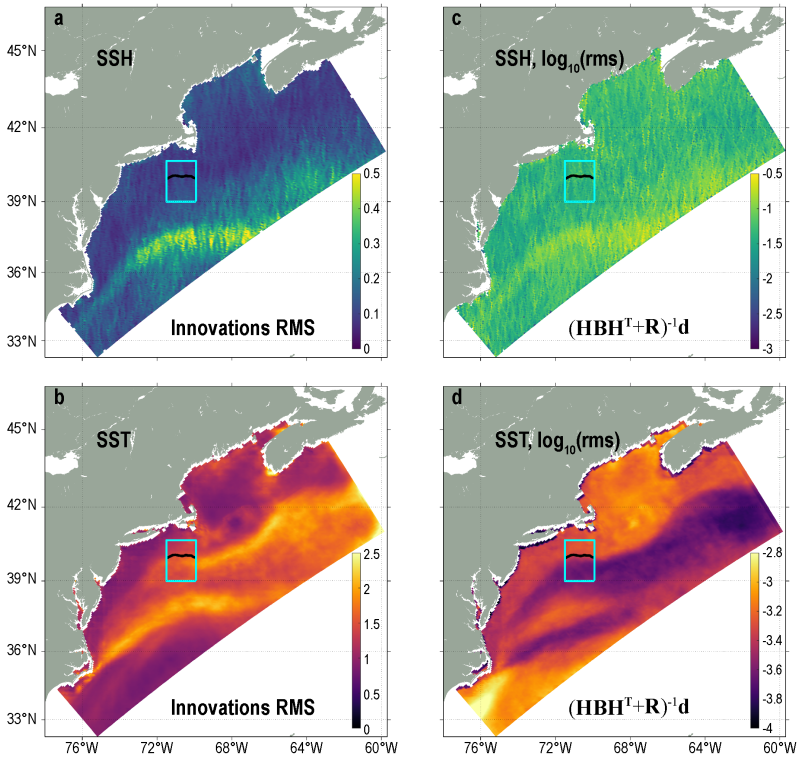

Figure 12: The RMS innovations **d** for (a) SSH and (b) SST averaged over all 4D-Var cycles (2014-2017). Also shown is $\log_{10}$ of the RMS of the product of the inverse stabilized representer matrix and the innovations for (c) SSH and (d) SST. The target section of the 200 m isobath for each metric is shown for reference crossing the Pioneer glider array (green box).

and 9d confirms that some of the regions of elevated impact are associated with large innovations in SST (such as those associated with the shelf-break and Gulf Stream fronts) and SSH (such as the Gulf Stream), in agreement with Fig. 8. However, Fig. 12c indicates that over the continental shelf $(\mathbf{HBH}^T + \mathbf{R})^{-1}$ accounts for much of the structure of the impact of SSH that is apparent in Fig. 9. On the other hand, for SST, Fig. 12d shows that $(\mathbf{HBH}^T + \mathbf{R})^{-1}$ generally down-weights the innovations in the frontal regions. The emergence of additional structure in Fig. 11f is associated with $\mathbf{HBM}^T(\partial I/\partial \mathbf{x})|_{\mathbf{x}^b}$, indicating that the impact structure for SST in Fig. 9 is also significantly controlled by the underlying circulation. Similar analyses to those presented in Fig. 12 for

the case of I_{uT} and I_{uS} (not shown) support these conclusions.

8. Observation Sensitivity and Borrowing Strength

The observation impact calculations of section 6 quantify the contribution of each observation to the metric increments resulting from data assimilation. However, these impacts do not necessarily indicate how the metrics themselves will change if particular observations are excluded from the 4D-Var analysis cycles. This information is provided by a complementary calculation which quantifies the *sensitivity* of a metric I to changes in the observation values or observing array (Moore et al., 2011c). In this case, the analysis equation (1) is expressed in a more general form as:

$$\mathbf{x}^a = \mathbf{x}^b + \mathcal{K}(\mathbf{d}) \tag{10}$$

where $\mathcal{K}(\mathbf{d})$ represents the *entire* data assimilation algorithm as a function of the innovation vector $\mathbf{d} = \mathbf{y}^o - H(\mathbf{x}^b)$. To first-order, a change $\delta\mathbf{y}^o$ in the observations will lead to a change in the analysis $\delta\mathbf{x}^a = (\partial\mathcal{K}/\partial\mathbf{y}^o)|_{\mathbf{x}^b}\delta\mathbf{y}^o$ where $(\partial\mathcal{K}/\partial\mathbf{y}^o)|_{\mathbf{x}^b}$ represents the tangent linearization of the entire 4D-Var data assimilation system about the background $\mathbf{x}^b$. At the end of the 1st outer-loop, the resulting change in any circulation metric to first-order is then given by:

$$\delta I_L \simeq \delta\mathbf{y}^{oT}(\partial\mathcal{K}/\partial\mathbf{y}^o)|^T_{\mathbf{x}^b} \sum_{i=i_s}^{i_e} \mathbf{M}_i^T \mathbf{h}_i \tag{11}$$

$$\delta I_Q \simeq \delta\mathbf{y}^{oT}(\partial\mathcal{K}/\partial\mathbf{y}^o)|^T_{\mathbf{x}^b} \sum_{i=i_s}^{i_e} \mathbf{M}_i^T (\mathbf{E}_i + \mathbf{E}_i^T)\mathbf{x}_i^b \tag{12}$$

for the general linear and quadratic classes of metric considered in section 6. In (11) and (12), $(\partial\mathcal{K}/\partial\mathbf{y}^o)|^T_{\mathbf{x}^b}$ represents the adjoint of the entire 4D-Var data assimilation system.

As shown by Moore et al. (2011b), choosing $\delta\mathbf{y}^o$ equal to -1 times the innovations associated with a particular subset of observations allows us to estimate the change that will occur in a metric if those observations are *excluded* from the 4D-Var analysis calculation. In particular, the change that occurs

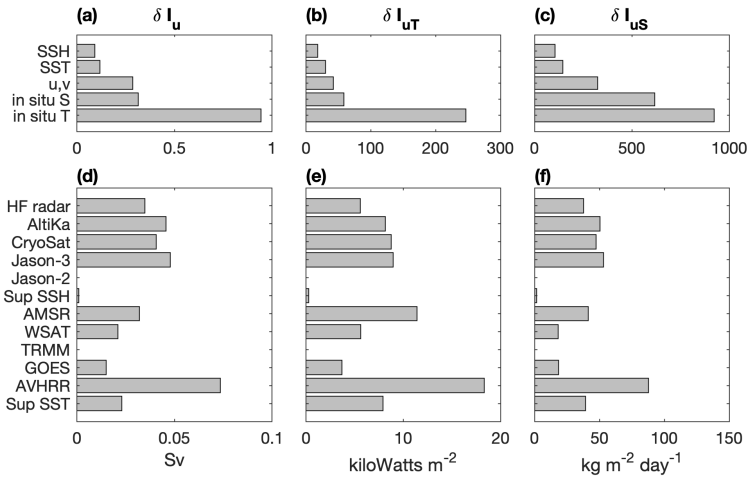

Figure 13: Histograms of the RMS change in each metric averaged over all 2017 4D-Var cycles when observations of different types and from different remote sensing platforms are excluded from the 4D-Var analyses, for (a,d) δI_u, (b,e) δI_{uT} and (c,f) δI_{uS}.

in I_L or I_Q due to eliminating any subset of observations from any 4D-Var cycle can be computed from a single integration of the adjoint of the 4D-Var algorithm. Figure 13 shows the RMS change in I_u, I_{uT}, and I_{uS} that results from independently eliminating all of the observations of the type indicated from all of the 4D-Var analysis cycles during 2017. Since the computational cost of the adjoint of the 4D-Var system is comparable to that of 4D-Var itself, we have limited these calculations to this single year when data coverage is good for all platforms. A histogram of the RMS impacts for 2017 alone (not shown) is very similar to Fig. 6 except that very few Jason-2 altimeter observations were used during 2017. This was due to Jason-2 having been superseded by the Jason-3 tandem mission, and our decision not to use data from the Jason-2 geodetic mission that began in July 2017 out of concern that the data errors might be increased by possible inaccuracies in the MSS and MDT off the reference mission ground-tracks. (We have subsequently regained confidence in these data and restored Jason-2 to the data flow for the MARACOOS model.)

A comparison of Fig. 13 with the RMS impacts of Fig. 6 reveals several noteworthy differences between the conclusions drawn from the impact calculations and the supporting sensitivity calculations. First, while the observation impact calculations of Fig. 6 indicate that both *in situ* temperature and salinity observations have the largest impact on the three transport metrics, and contribute approximately equally to the metric increments, Fig. 13 reveals that the elimination of all the *in situ* temperature observations from each analysis cycle yields the largest change in each metric. Second, while Fig. 6 shows that *in situ* velocity observations and HF radar observations have a relatively low impact on each metric, Fig. 13 indicates that elimination of these data from each 4D-Var cycle leads to a relatively large change in all metrics, comparable in fact to eliminating both SST and SSH. Third, Fig. 13 shows that excluding satellite SSH and SST observations leads to smaller changes in the metric increments than might be anticipated from the impact calculations of Fig. 6.

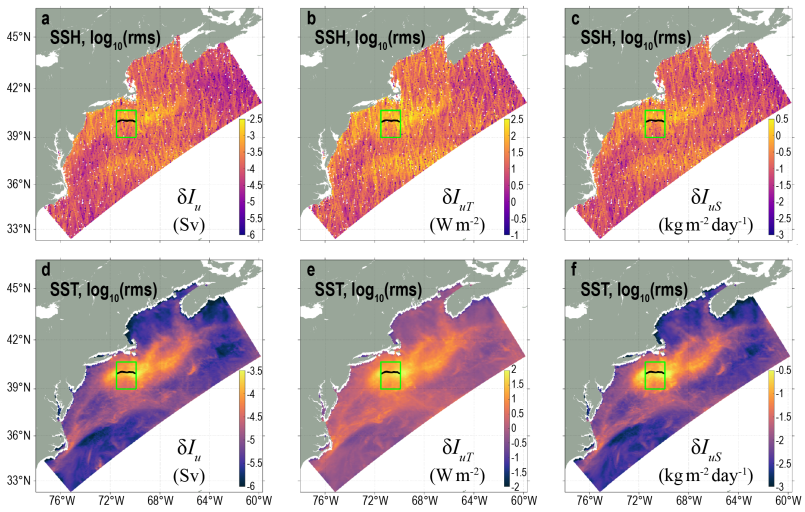

Figure 14: The spatial distribution of the $\log_{10}$ of the RMS change that results in each index when observations falling within each model are eliminated from all 2017 4D-Var analyses. Specifically for altimeter observations and (a) δI_u , (b) δI_{uT} and (c) δI_{uS}, and satellite SST observations for (d) δI_u, (e) δI_{uT} and (f) δI_{uS}. In each case the RMS is shown at the end of the 1st outer-loop. The location of the 200 m isobath target section crossing the Pioneer glider array (green box) is indicated.

The sensitivity of the metric increments to exclusion of SST and SSH observations is further explored in Fig. 14, which shows the RMS change in each metric when all observations that fall within a given grid cell are eliminated from the 2017 4D-Var cycles. Fig. 14 shows that exclusion of observations near the target section generally yields the largest change in each metric. For SSH, this is in stark contrast to Fig. 9, which indicates that it is typically SSH observations remote from the target section that have the most significant impact on each metric. For SST, Fig. 9 indicates that while observations in the vicinity of the target section have an impact, there are significant impacts from remote observations as well, in contrast to Fig. 14.

We hypothesize that the seemingly contradictory nature of the observation impact and observation sensitivity calculations can be understood in terms of the concept of *borrowing strength* that arises in the field of statistics. The idea of borrowing strength was first introduced in the 1960s and 70s by the mathematician John Tukey at Princeton University in connection with election night analysis and forecasting of U.S. Congressional seats (see Brillinger, 2002), and, since then, has enjoyed tremendous popularity in the fields of empirical Bayesian methods and shrinkage estimation for ill-posed problems. The idea is valid in the current context since 4D-Var represents a form of penalized least-squares in which the ill-posed nature of the ocean state estimation problem is regularized using the background error covariance information in $\mathbf{B}$ (e.g., see Bennett, 2002). The basic idea behind borrowing strength is that while some observations, such as those from HF radars, do not have a significant direct influence on a chosen metric I as quantified by the observation impact calculations, they can nonetheless provide information about the ocean state which corroborates that obtained by other observing platforms, thus enhancing the utility of the latter. In this way, the HF radar observations can have a sizable *indirect* influence on I if they are subsequently excluded from the analysis, since the corroborating information that they provide about the ocean state in support of other observing platforms is lost.

To illustrate this idea further, Table 4 shows the ratio of the RMS *sensitiv-*

ity per datum to the RMS *impact* per datum for all three metrics. This ratio is a measure of the average change that actually occurs in each metric when an observation is excluded from all 4D-Var analysis and that which might be expected to occur based on the observation impact alone. Thus, a ratio of one indicates that the sensitivity and impact calculations predict, on average, the same change in a given metric if an observation is excluded from each 4D-Var cycle. On the other hand, a ratio greater than one indicates that the actual change in an index will be larger than expected based on the observation impact calculations. Thus the departure of the ratios in Table 4 from a value of one can be viewed as an indicator of the level of borrowing strength. Inspection of Table 4 reveals that, in fact, all observation platforms appear to be borrowing strength from each other to a different degree. The observations that appear to lend most support to the other elements of the observing system are measurements of surface currents by HF radar observations. Conversely, altimetry measurements do not appear to add much additional support to the observing system than is suggested by the observation impact analyses. Since ratios in Table 4 for all observation platforms are greater than one, this indicates that the observing array is stronger in its entirety than the consideration of each platform separately would suggest (*i.e.* each observation platform benefits from the collective presence of all other platforms).

Returning to the HF radar and velocity observations, Fig. 6 shows that these data have relatively little direct impact on the increments in the three transport metrics. However, they do in fact contribute useful guidance to the 4D-Var analyses that corroborates information from other platforms as evidenced in Fig. 13. In fact, Table 4 shows that these data are an order of magnitude more impactful than the observation impact analyses suggest. It is useful to speculate in what ways data from other observing platforms can borrow strength from the HF radar and in situ velocity observations. For example, the surface current estimates from HF radar observations comprise contributions from the wind-driven Ekman and pressure-driven flows. Satellite altimeter observations, on the other hand, can only detect the signature of the pressure-driven flow, in-

	I_u	I_{uT}	I_{uS}
SST	4.9	6.0	5.9
SSH (all)	1.6	1.4	1.3
in situ T	3.3	3.0	3.3
in situ S	3.2	3.2	4.0
in situ u,v	6.7	5.4	6.7
HF radar	10.4	9.4	10.4

Table 4: The ratio of RMS sensitivity per datum to the RMS impact per datum for each transport metric during 2017. This ratio is a measure of the average change that actually occurs in each metric when an observation is excluded from the 4D-Var analysis and that which might be expected to occur based on the observation impact alone. The repeated assimilation of SSH observations described in section 4.1 has been accounted for here.

formation that will be corroborated by the HF radar data. Furthermore, much of the energy in the circulation will be in the form of potential energy (except at the smallest scales), and, as such, the temperature and salinity observations that provide direct information about the density field will generally be more effective at recovering the circulation than velocity observations. Nonetheless, the velocity observations will contain the, albeit small, kinetic energy signatures associated with coherent circulation features that are resolved by remote sensing and hydrographic observations, and the velocity observations will corroborate the information from the latter observing platforms. Thus, in this way, remote sensing and *in situ* instruments can be said to be borrowing strength from the velocity observations. Consequently, if the HF radar and other velocity observations are excluded from the 4D-Var analysis, the corroborating information that they provide will be lost, and the observation sensitivity calculations indicate that this can have a far greater impact on the analyses than observation impact calculations alone suggest.

Similarly, the differences between the spatial distributions of the SST and SSH impacts in Fig. 9 and the corresponding sensitivities of Fig. 14 can also be interpreted in terms of borrowing strength. The high sensitivity of the cross-

shelf transport metrics to observations in the vicinity of the target regions in Fig. 14 indicates that the *in situ* observations from the Pioneer array are benefiting (*i.e.*, borrowing strength) from the corroborating information provided by satellite remote sensing data in this region. Therefore, while not shown here, the Pioneer array *in situ* observations have the largest impact locally on the cross-shelf metrics, and the loss of remote sensing observations from this region will significantly impact each metric because the corroborating information the remote sensing platforms provide to the *in situ* observations will be lost.

The observation impact and observation sensitivity calculations therefore provide complementary information about the direct and indirect impact of each observation on a chosen metric, and furnish direct evidence for the synergy that exists within the 4D-Var system between observations from different platforms.

9. Summary and Conclusions

The impact of observations from *in situ* sensors and satellite and terrestrial remote sensing platforms on cross-shelf transport in the vicinity of the MAB front derived from 4D-Var estimates of the ocean state has been quantified using a rigorous method that is currently employed in numerical weather prediction at several operational centers. Despite being relatively small in number, the data that have the largest impact are *in situ* measurements of temperature and salinity since these provide direct information about the 3-dimensional structure of the circulation. However, satellite remote sensing observations of SSH and SST collectively also exert a significant influence on 4D-Var circulation estimates. While the impact of an individual SST observation is relatively small compared to that of an *in situ* datum (cf Table 3), the large volume of SST data yields a sizeable aggregate impact. Along track altimetry observations, on the other hand, are generally the least abundant data stream (cf Fig. 2), but the impact per observation is more than an order of magnitude larger than that of SST. Therefore, the aggregate impact of altimetry observations is on par with that of SST. While the overall impact of HF radar observations appears to be quite

modest, these data do, in fact, contribute important information from which other data streams borrow strength as discussed below.

The impact information can be viewed as one quantitative measure of the value of data from different components of the ocean observing systems. In addition though, the observation impacts also provide useful diagnostics for assessing the performance characteristics of the data assimilation system, and yield a quantitative measure of the utility of observations from different observing platforms. This information can be used for tuning the data assimilation system with the aim of extracting additional information from under-utilized observing platforms.

The impact of an observation within the 4D-Var system will depend on several factors that include the extent to which information is dynamically interpolated in space and time by the background circulation (via the tangent linear and adjoint models), and the covariance between the observation locations of errors and uncertainties in the background and the observations. The geographic distributions of the impacts were found to be remarkably robust across several metrics, across different platforms, and through time.

From an observing system design perspective, these patterns could be taken as guidance on regions where acquiring sustained observations will contribute the most information toward constraining cross-shelf exchange estimates, noting that these regions are not always local. However, it must be emphasized that all the specifics of the results presented here are for the particular flow indices, I_u, I_{uT} and I_{uS}, that we defined *and* the geographic section at which they are evaluated. Extension to other dynamic or application-specific indices, and locations, is readily pursued by defining new metrics following (3) - (5), deriving $(\partial I/\partial \mathbf{x})|_{\mathbf{x}^b}$, and applying the methodology outlined above.

A step-by-step analysis of the observation impact procedure revealed that the robust geographic distributions of impact are the result of the combined influence of the space-time dynamic interpolation of 4D-Var and the prior error covariance information. Therefore, unravelling the flow of information through the data assimilation system is not straightforward. Clearly though the space-

time dynamic interpolation relies on familiar dynamical processes such as advection by ocean currents and wave propagation, and it is useful to think of the *information horizons* associated with these processes along the lines discussed by Moore et al. (2015). The primary circulation features that control the extent of the advection information horizons are the shelf-break jet ($\sim 0.1\text{m s}^{-1}$) and the Gulf Stream ($\sim 2\text{m s}^{-1}$). Over a 3-day assimilation cycle, information can therefore be advected ~ 25 km($\sim 0.3^{\circ}$) by the the shelf-break jet and ~ 500 km($\sim 6^{\circ}$) by the Gulf Stream. The information horizons associated with wave propagation will depend on the propagation speed of baroclinic and barotropic waves. The fastest internal waves will travel at the first baroclinic mode phase speed, which in this region is $\sim 2\text{ m s}^{-1}$ (Chelton et al., 1998). Therefore the extent of the baroclinic wave information horizon will be similar to that due to advection by the Gulf Stream. The very rapid phase speed of the barotropic mode ensures that its information horizon will encompass the entire model domain. With these factors in mind, the geographic patterns of impact evident in Fig. 9 are perfectly plausible based on the dynamic propagation of information alone.

As noted by Trémolet (2008), the impact of observations at *any* stage of the 4D-Var calculation can be quantified by archiving appropriate information from the data assimilation system. For the ROMS 4D-Var platform this idea raises some interesting possibilities since, in principle, the contribution of each term in the governing equations to the observation impact calculations can be quantified. Therefore, a future study aimed at quantifying the role of advection, diffusion, pressure gradient effects, etc, on the flow of information through the 4D-Var system is envisaged.

The synergy between observations from different platforms was identified in the present study by comparing the outcome of the observation impacts calculations with the outcome of observation sensitivity calculations. This comparison uncovers the extent to which observations from different platforms are able to borrow strength from each other. In this way, observations from seemingly low impact platforms, such as HF radar surface current estimates, can provide in-

formation about the state of the ocean which corroborates that available from other platforms, such as satellite systems. This inherent behavior of data assimilation systems behooves us to pay close attention to the utility and attributes of *all* observing platforms, since the elimination of a seemingly low impact data stream could significantly degrade the performance of an analysis-forecast system. These types of analyses also pave the way for optimizing the design of observing systems.

Acknowledgements

This work was supported by grants from the National Science Foundation (OCE-1459665 and OCE-1459646), NASA (NNX17AH58G) and NOAA (NA16NOS0120020).

References

Banerjee, S., Carlin, B., Gelfand, A., 2004. Hierarchical modeling and analysis for spatial data. volume 101. Chapman and Hall/CRC Press, Boca Raton, FL.

Bennett, A.F., 2002. Inverse Modeling of the Ocean and Atmosphere. Cambridge University Press, Cambridge, U.K.

Boyer, T., Antonov, J., Baranova, O., Garcia, H., Johnson, D., Locarnini, R., Mishonov, A., OBrien, T., Seidov, D., Smolyar, I., Zweng, M., 2009. World ocean database 2009, s. levitus, ed. NOAA Atlas NESDIS 66 , 216pp.

Brillinger, D., 2002. John w. tukey: His life and professional contributions. Annals of Statistics 30, 1535–1575.

Chelton, D., DeSzoeke, R., Schlax, M., Naggar, K.E., Siwertz, N., 1998. Geographical variability of the first baroclinic rossby radius of deformation. J. Phys. Oceanogr. 28, 433–460.

Chen, K., He, R., 2010. Numerical investigation of the Middle Atlantic Bight shelfbreak frontal circulation using a high-resolution ocean hindcast model. J. Phys. Oceanogr. 40, 948–964.

Courtier, P., Thépaut, J.N., Hollingsworth, A., 1994. A strategy for operational implemenation of 4D-Var using an incremental approach. Quarterly Journal of the Royal Meteorological Society 120, 1367–1388.

Daley, R., 1991. Atmospheric Data Analysis. Cambridge University Press, Cambridge, U.K.

Drévillon, M., Bourdallé-Badie, R., Derval, C., Lellouche, J.M., Rémy, E., Tranchant, B., Benkiran, M., Greiner, E., Guinehut, S., Verbrugge, N., Garric, G., Testut, C.E., Laborie, M., Nouel, L., Bahurel, P., Bricaud, C., Crosnier, L., Dombrowsky, E., Durand, E., Ferry, N., Hernandez, F., Galloudec, O.L., Messal, F., Parent, L., 2008. The GODAE/Mercator-Ocean global ocean forecasting system: results, applications and prospects. Journal of Operational Oceanography 1, 51–57.

Fairall, C., Bradley, E., Hare, J., Grachev, A., Ebson, J., Young, G., 2003. Bulk parameterization of airsea fluxes: Updates and verification for the COARE algorithm. J. Climate 16, 571–591.

Fleming, N., 2016. Seasonal and spatial variability in temperature, salinity and circulation of the Middle Atlantic Bight. Ph.D. thesis. Rutgers University, New Brunswick, N.J.

Garvine, R.W., Wong, K.C., Gawarkiewicz, G., 1989. Quantitative properties of shelfbreak eddies. J. Geophys. Res. 94, 14475–14483.

Gawarkiewicz, G., Todd, R., Zhang, W., Partida, J., Gangopadhyay, A., Monim, M.U.H., Fratantoni, P., Malik, A.M., Dent, M., 2018. The changing nature of shelf-break exchange revealed by the OOI Pioneer Array. Oceanography 31, 60–90.

Gürol, S., Weaver, A., Moore, A., Piacentini, A., Arango, H., Gratton, S., 2014. B-preconditioned minimization algorithms for variational data assimilation with the dual formulation. Quarterly Journal of the Royal Meteorological Society 140, 539–556.

Gratton, S., Tshimanga, J., 2009. An observation-space formulation of variational assimilation using a restricted preconditioned conjugate gradient algorithm. Quarterly Journal of the Royal Meteorological Society 135, 1573–1585.

Ide, K., Courtier, P., Ghil, M., Lorenc, A., 1997. Unified notation for data assimilation: Operational, sequential and variational. J. Meteorol. Soc. Japan 75, 181–189.

Järvinen, H., Undén, P., 1997. Observation screen and quality control in the EMCWF 3D Var data assimilation system. ECMWF Tech. Memo 236, 34.

Langland, R., Baker, N., 2004. Estimation of observation impact using the NRL atmospheric variational data assimilation adjoint system. Tellus 56A, 109–201.

Lawless, A., Gratton, S., Nichols, N., 2005. Approximate iterative methods for variational data assimilation. Int. J. Numer. Meth. Fl. 1, 1–6.

Levin, J., Wilkin, J., Fleming, N., Zavala-Garay, J., 2018. Mean circulation of the mid-atlantic bight from a climatological data assimilative model. Ocean Modelling 128, 1–14.

Linder, C., Gawarkiewicz, G., 1998. A climatology of the shelf break front in the Middle Atlantic Bight. J. Geophys. Res. 102, 18405–18423.

Moore, A., Arango, H., Broquet, G., Edwards, C., Veneziani, M., Foley, B.P.D., Doyle, J., Costa, D., Robinson, P., 2011a. The Regional Ocean Modeling System (ROMS) 4-dimensional variational data assimilation systems. Part II: Performance and application to the California Current System. Progress in Oceanography 91, 50–73.

Moore, A., Arango, H., Broquet, G., Edwards, C., Veneziani, M., Foley, B.P.D., Doyle, J., Costa, D., Robinson, P., 2011b. The Regional Ocean Modeling System (ROMS) 4-dimensional variational data assimilation systems. Part III: Observation impact and observation sensitivity in the California Current System. Progress Oceanography 91, 74–94.

Moore, A., Arango, H., Broquet, G., Powell, B., Zavala-Garay, J., Weaver, A., 2011c. The Regional Ocean Modeling System (ROMS) 4-dimensional variational data assimilation systems. Part I: System overview and formulation. Progress in Oceanography 91, 34–49.

Moore, A., Edwards, C., Fiechter, J., Drake, P., Arango, H., Neveu, E., Gürol, S., Weaver, A., 2013. A prototype for an operational regional ocean data assimilation system. In "Data Assimilation for Atmospheric, Oceanic and Hydrological Applications, Vol. II," Liang Xu and Seon Park, Eds. Springer, Berlin Heidelberg, Chapter 14 , 345–366.

Moore, A., Edwards, C., Fiechter, J., Jacox, M., 2015. Observing system impacts on estimates of california current transport. In "Coastal Ocean Observing Systems: Advances and Synthesis," Yonggang Liu, Heather Kekering and Robert Weisberg, Eds., Academic Press, Chapter 19 , 351–369.

Moore, A., Jacox, M., Crawford, W., Laughlin, B., Edwards, C., Fiechter, J., 2017. The impact of the ocean observing system on estimates of the California Current circulation spanning three decades. Progress in Oceanography 156, 41–60.

Mukai, A.Y., Westerink, J.J., Luettich, R.A., D. Mark, ., 2002. A tidal constituent database for the western North Atlantic, Gulf of Mexico and Caribbean Sea. Tech. Rep. ERDC/CHL TR-02-24 , 196pp.

Ray, R., 2013. Precise comparisons of bottom-pressure and altimetric ocean tides. J. Geophys. Res. 118, 4570–4584.

Schaeffer, P., Pujol, M.I., Faugere, Y., Picot, N., Guillot, A., 2016. New mean sea surface CNES-CLS 2015 focusing on the use of geodetic missions of CryoSat-2 and Jason-1. Presentation at Living Planet Symposium, Prague, Czech Republic, 9-13 May 2016 .

Talagrand, O., Courtier, P., 1987. Variational assimilation of meteorological observations with the adjoint vorticity equation. I: Theory. Quarterly Journal of the Royal Meteorological Society 113, 1311–1328.

Trémolet, Y., 2008. Computation of observation sensitivity and observation impact in incremental variational data assimilation. Tellus 60A, 964–978.

Weaver, A., Courtier, P., 2001. Correlation modelling on the sphere using a generalized diffusion equation. Quarterly Journal of the Royal Meteorological Society 127, 1815–1846.

Wilkin, J., Levin, J., Lopez, A., Hunter, E., Zavala-Garay, J., Arango, H., 2018. A Coastal Ocean Forecast System for U.S. Mid-Atlantic Bight and Gulf of Maine. In *New Frontiers in Operational Oceanography*, E.P. Chassignet, A. Pascual, J. Tintoré and J. Verron, Eds., GODAE OceanView Chapter 21, 593–624. doi:10.17125/gov2018.ch21.

Zavala-Garay, J., Wilkin, J., Levin, J., 2014. Data assimilation in coastal oceanography: IS4DVAR in the Regional Ocean Modeling System (ROMS). Advanced Data Assimilation for Geosciences: Lecture Notes of the Les Houches School of Physics: Special issue June 2012, E. Blayo, M. Bocquet, E. Cosme, and L. Cugliandolo, Eds , 555–576.

Zhang, W., Gawarkiewicz, G., 2015. Dynamics of the direct intrusion of Gulf Stream ring water onto the Mid-Atlantic Bight shelf. Geophysical Research Letters 42, 7687–7695.