Continental Shelf Research 29 (2009) 1711–1722

Contents lists available at ScienceDirect

Continental Shelf Research

journal homepage: www.elsevier.com/locate/csr

Mapping radar-derived sea surface currents with a variational method

Max Yaremchuk [a,*], Alexei Sentchev [b]

[a] Department of Physics, University of New Orleans, LA 70143, USA
[b] Laboratorie d'Océanologie et Géosciences, Université du Littoral, Wimereux, France

ARTICLE INFO

Article history:
Received 22 December 2008
Received in revised form
13 May 2009
Accepted 28 May 2009
Available online 18 June 2009

Keywords:
High-frequency radars
Interpolation
Variational method

ABSTRACT

High-frequency radars measure projections of surface velocity vectors on the directions of the radar beams. A variational method for reconstruction of the 2d velocity field from such observations is proposed. The interpolation problem is regularized by penalizing high-frequency variability of the surface vorticity and divergence fields. Twin-data experiments are used to assess the method's skill and compare it with two well-known approaches to HFR data processing: conventional local interpolation and more sophisticated non-local scheme known as open-boundary modal analysis (OMA). It is shown that the variational method and OMA have a significant advantage over local interpolation because of their ability to reconstruct the velocity field within the gaps in data coverage, near the coastlines and in the areas covered only by one radar. Compared to OMA, the proposed variational method appears to be more flexible in processing gappy observations and more accurate at noise levels below 30%.

1. Introduction

The technology of monitoring near-coastal currents by high frequency radars (HFRs) has been rapidly developing in the past decade. HFR observations were extensively used to study near-shore circulation under a large variety of environmental conditions (e.g., Prandle, 1993; Graber et al., 1996; Hisaki et al., 2001; Breivik and Sætra, 2001; Sentchev and Yaremchuk, 2007; Shay et al., 2007; Chavanne et al., 2007). At present, the HFR surface current mapping technology is capable to provide surface velocity maps in an area up to 200 km offshore, with space–time resolutions of 0.2–1 h/1–10 km depending on the particular system and practical application.

An important question in dealing with HFR data is the problem of retrieving the 2d velocity vector maps from the velocity components measured along the beams. The most commonly used technique (e.g., Shay et al., 2007) is based upon local interpolation (LI) of the radial data. It takes into account both measurement errors and the effect of geometric dilution of precision (Chapman et al., 1997) associated with the angle between the intersecting beams in a given grid cell.

Recently, Park et al. (2007) proposed a more sophisticated algorithm for reconstruction of the stream function ψ and velocity potential φ associated with the surface velocity field. The algorithm is based on minimization of a cost function in the ψ/φ space. The problem is regularized by attracting the interpolation pattern to a prior background state in a metric associated with a user-specified covariance function. The latter has to be simple enough to be inverted in realistic applications.

Another algorithm which has been under development in recent years (Lekien and Coulliette, 2004; Kaplan and Lekien, 2007) in the HFR community is the open-boundary modal analysis (OMA). This technique can be viewed as an extension of the normal mode analysis (Eremeev et al., 1992; Lipphardt et al., 2000) which, similarly to Park et al. (2007), employs decomposition of a 2d vector field into divergent and rotational components. Apart from the capability to avoid explicit specification of the poorly known error covariance, the OMA technique automatically takes into account the kinematic constraints imposed on the velocity field by the coastlines.

The OMA is based on expansion of ψ and φ in the eigenfunctions of the Laplacian operators (interior modes) and solutions to specific boundary value problems (boundary modes). Regularization is achieved by an appropriate spectrum cutoff and penalizing the magnitudes of the remaining low-wavenumber spectral components. Such an approach may lack flexibility in representing localized small-scale features as well as flow structures near open boundaries, where the rotational component of the velocity field is subject to the zero normal flow constraint. Besides, certain difficulties may arise when dealing with gappy data, especially when the horizontal size of a gap is larger than the minimal length scale resolved (Kaplan and Lekien, 2007).

In this paper we present a 2d variational (2dVar) interpolation method designed for HFR data processing and compare its performance with the LI and OMA techniques. In contrast to the method of Park et al. (2007), the proposed algorithm does not require specification of a background velocity field and its

* Corresponding author. NRL Code 7321, Bldg. 1009, Stennis Space Center, MS 39529, USA.
E-mail address: myaremch@uno.edu (M. Yaremchuk).

0278-4343/$ - see front matter © 2009 Elsevier Ltd. All rights reserved.
doi:10.1016/j.csr.2009.05.016

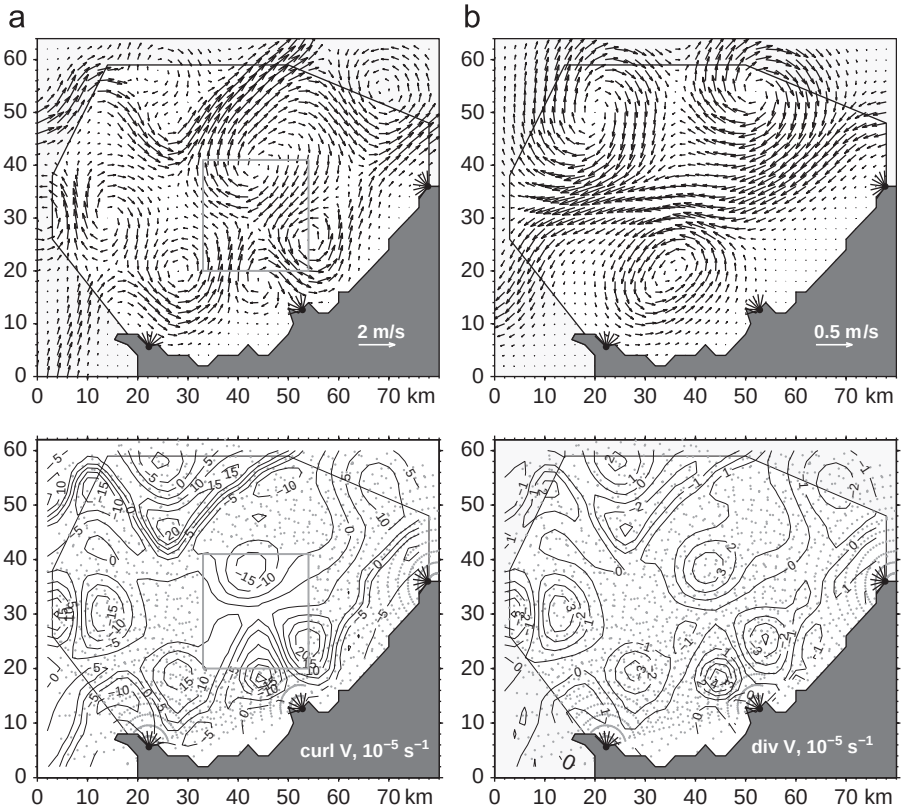

Fig. 1. Setting of the numerical experiments: upper panels show two velocity fields used in twin-data experiments. The first field (A) has nine eddy structures superimposed on two jets: one following the coastline and another located farther offshore. The second one (B) has three large eddies and a broad jet. Lower panels show vorticity (left) and divergence (right) fields, corresponding to A. Three coastal radars sample the radial velocities along beam directions binned at 2 km radial and 5° azimuthal resolution. Sampling points are shown by gray dots. Gray rectangle in the lower left panel envelops simulated gap in the HFR data. Radar positions are shown by black dots. The interpolation domain and the coastline are similar to the ones used in KL07 for OMA processing of HFR observations in the Bodega Bay. The coordinate system is rotated clockwise (north is on the right). Contour intervals for vorticity and divergence are $5\times10^{-5}\,s^{-1}$ and $10^{-5}\,s^{-1}$, respectively.

covariance structure that may be lacking in many applications. Similar to OMA, the 2dVar method is non-local (the result of interpolation at a given grid point depends on all the observed radial velocities) and kinematically constrained (the interpolated velocity field $\mathbf{v}$ is subject to constraints imposed by the coastline).

A distinct feature of the 2dVar approach is that the interpolating functions (eigenvectors of the Hessian matrix) depend both on the shape of the coastline and on the spatial distribution of the data points. The latter feature (lacking in OMA) is especially useful in practice, because HFR data often have gaps in spatial coverage, and the issue of filling those gaps is important.

Another useful property of the proposed 2dVar method is that, similar to LI, it operates in the velocity space, and thus requires less sophisticated operators for projection of the unknown gridded velocities on the radial components of the current speed. This property provides better conditioning of the Hessian matrix, faster convergence of the minimization algorithm and improved computational efficiency.

This paper is organized as follows: in the next section we describe the algorithm by specifying the cost function and kinematic constraints, outline its numerical implementation and the error analysis scheme. Section 3 is devoted to comparison of the method with LI and OMA techniques in the framework of numerical experiments with simulated data in a realistic domain. In Section 4 we compare the real HFR data interpolation by the LI, OMA and 2dVar methods. Discussion and conclusions complete this paper.

2. Methodology

Consider an oceanic domain Ω partly bounded by the coastline $\partial\Omega_0$ where HFRs are located (Fig. 1). Projections v^* of the surface velocity field $\mathbf{v}^t(x,y)$ on the radar beam directions $\mathbf{r}$ are observed at a discrete set of points $\mathbf{x}_k, k=1,\ldots,K$ located along the beams. Our goal is to obtain an estimate $\mathbf{v}$ of $\mathbf{v}^t$ given v_k^* and their observation errors $\sigma(v_k^*)$.

2.1. Cost function

Following the standard methodology of the variational interpolation the "best" estimate $\mathbf{v}$ is determined by maximizing its likelihood expressed in terms of the Gaussian probability density function $\mathscr{P}(\mathbf{v})\sim\exp[-J(\mathbf{v})]$ (e.g., Thacker, 1989). The argument of the exponent (the cost function) is quadratic in $\mathbf{v}$ and usually consists of two terms: $J=J_d+J_r$.

The first term J_d measures the "distance" between observations v^* and the corresponding components of the unknown field $\mathbf{v}$ at the observation points:

$$J_d=\frac{1}{2}\left[\sum_{k=1}^{K}\sigma^{-2}(v_k^*)[(\mathbf{v}\cdot\mathbf{r}_k)-v_k^*]^2\right] \tag{1}$$

Because of Gaussianity the "distance" is quadratic in $\mathbf{v}$ and scaled in the observation points by the corresponding error variances. Although the form of J_d implicitly assumes that observation errors are uncorrelated, a more general formulation can also be

considered by specifying the full $K \times K$ inverse error covariance matrix. In practice, error covariances between the HFR radar observations are known with much less accuracy than their variances, and we will further consider the case of the diagonal observational error covariance matrix.

The second term J_r is introduced to regularize the interpolation problem: the unknown continuous field $\mathbf{v}$ has many more degrees of freedom than the number of observations K. Here we follow a standard regularization approach (e.g., Thacker, 1988) and penalize higher spatial derivatives of $\mathbf{v}$:

$$J_r = \frac{K}{2A}\int_\Omega [W^d(x,y)(\Delta \mathrm{div}\mathbf{v})^2 + W^c(x,y)(\Delta \mathrm{curl}\mathbf{v})^2]\,dx\,dy \tag{2}$$

where $\mathrm{div}\mathbf{v} = \partial_x u + \partial_y v$, $\mathrm{curl}\mathbf{v} = \partial_x v - \partial_y u$ are the divergence and vorticity, respectively, $\Delta = \partial_{xx} + \partial_{yy}$ is the Laplacian operator and A is the area of interpolation domain bounded in Fig. 1 by a thin line.

Conceptually, it is convenient to represent $\Delta\mathrm{div}\mathbf{v}$ and $\Delta\mathrm{curl}\mathbf{v}$ as additional Gaussian δ-correlated fields, whose "observed" values are zeroes and error variances are $W_d^{-1}(x,y)$ and $W_c^{-1}(x,y)$, respectively (Thacker, 1988). By introducing such bogus observations in every spatial location, we effectively make the total number of data points larger than the number of unknowns, making the interpolation problem formally well posed. Note that since the Laplacian acts as a high-pass filter, the assumption on zero spatial correlations between the bogus data appears to be much more realistic as compared to the case if we penalized the magnitudes of $\mathrm{div}\mathbf{v}$ and $\mathrm{curl}\mathbf{v}$ (i.e. "observed" their zero values).

In contrast to $\sigma^{-2}(v_k^*)$ that can be estimated experimentally with some confidence, the functions of W^d and W^c are largely unknown and should be obtained from *ad hoc* considerations. In the present study we set their spatial variability to be inversely proportional to the local density $\rho(x,y)$ of HFR observation points: $W^{c,d} = \alpha^{c,d}\rho^{-1}$. By introducing this assumption we expect to resolve finer structures of the velocity field in the regions of higher data density, and therefore relax the smoothness constraint. The definition of the scalar parameters α^c and α^d is given in Section 3.1.

After we specified the probability density function on the space $\mathscr{V}$ of all possible configurations of $\mathbf{v}(x,y)$, the "best" interpolation pattern is sought by maximizing the probability $\exp(-J)$, or, equivalently, by minimizing J with respect to $\mathbf{v}$. Technically, the minimization can be done in a subspace $\mathscr{V}_0 \subset \mathscr{V}$ constrained by additional conditions imposed on $\mathbf{v}$ by dynamics or some other reliable source of information. In the present study we utilize the simplest possible constraint by specifying the condition $\mathbf{v}(\partial\Omega_0) = 0$ at the coastline, i.e. the regular-grid velocity field is defined as the one minimizing J under the condition $\mathbf{v}(\partial\Omega_0) = 0$:

$$J_d + J_r \to \min_{\mathbf{v}\in\mathscr{V}_0}$$

2.2. Numerical implementation

The interpolated velocity $\mathbf{v} = \{u, v\}$ is discretized on a regular grid with a constant step δx in both directions. For a typical HFR network the domain size is always much smaller than the Earth radius, so we adopt flat geometry, assuming that spherical corrections are much smaller than uncertainties introduced by measurement errors and by the lack of knowledge of their covariance structure. First-order differential operators are approximated by central differences, whereas the discrete Laplacian of the scalar fields $\mathrm{div}\mathbf{v}$ and $\mathrm{curl}\mathbf{v}$ in Eq. (2) is defined by a standard 5-point stencil:

$$\tilde{\Delta}\varphi_{i,j} = [\varphi_{i,j+1} + \varphi_{i,j-1} + \varphi_{i+1,j} + \varphi_{i-1,j} - 4\varphi_{i,j}]/\delta x^2,$$

where φ stands for either divergence or vorticity and indices i,j enumerate regular grid points in x and y directions, respectively. Discretization converts J into a quadratic form defined on a finite-dimensional space of the grid point values of u and v:

$$\tilde{J} = \frac{1}{2}\left\{\sum_k \sigma^{-2}(v_k^*)[(\hat{P}_k^{ij}\mathbf{v}_{ij})\cdot\mathbf{r}_k - v_k^*]^2 + \frac{K}{N}\sum_{i,j}[W_{ij}^d(\tilde{\Delta}\tilde{\mathrm{div}}\mathbf{v}_{ij})^2 + W_{ij}^c(\tilde{\Delta}\tilde{\mathrm{curl}}\mathbf{v}_{ij})^2]\right\} \tag{3}$$

Here N is the number of regular-grid points in Ω and $\hat{P}_k^{ij}$ denote linear interpolation operators, which project the values of $\mathbf{v}_{ij}$ onto the HFR observation points $\mathbf{x}_k$ from the apexes of a grid cell, containing $\mathbf{x}_k$.

Constrained minimization of J is performed using the quasi-Newtonian descent algorithm of Gilbert and Lemarechal (1989). In addition to the computation of $\tilde{J}$ this algorithm also requires a procedure for computation of its derivatives $\partial\tilde{J}/\partial u_{ij}$ and $\partial\tilde{J}/\partial v_{ij}$. This procedure was obtained by analytical differentiation of the code for the computation of $\tilde{J}$.

2.3. Error estimation

Keeping in mind the probabilistic interpretation of $\tilde{J}$, we treat Eq. (3) as the argument of the Gaussian exponent. Under this assumption, the inverse error covariance matrix C^{-1} for gridded velocity $\mathbf{v}_{ij}$ is the matrix H of second derivatives of $\tilde{J}$ with respect to u_{ij} and v_{ij} (the Hessian matrix of the problem). Technically, the computation of H is implemented by taking the finite differences of $\partial\tilde{J}/\partial u_{ij}$ and $\partial\tilde{J}/\partial v_{ij}$ in $\mathscr{V}_0$. Because the gradient of $\tilde{J}$ is the linear function of u_{ij} and v_{ij}, the numerical approximation error of this procedure is zero.

It is easy to note that the matrix H may not necessarily have the full rank and, therefore, cannot be inverted to obtain the error covariance matrix C. This is a consequence of the fact that some of the velocity patterns cannot be observed by the HFR system. As a simple example consider HFR observations along a single beam: adding a velocity field orthogonal to that beam with vorticity and divergence linearly dependent on x and y does not change the value of the cost function. In realistic situations more complicated patterns may exist, which depend on the geometry of the domain, radar positions, sampling discretization along the beams, and the type of regularization.

Degeneracy of H implies that the result of interpolation depends on the velocity field $\mathbf{v}_0$ used as a "first guess" at the start of the minimization procedure: if unobserved patterns are present in $\mathbf{v}_0$, they will remain there unaffected. Consequently, in order to remove these unobserved fields from the interpolation pattern one has to start minimization from $\mathbf{v}_0 = 0$, as it is done in the presented algorithm. To remove the unobserved subspace from error analysis we perform eigenvalue decomposition of H with respect to a standard dot product in the velocity space and discard eigenvectors with zero eigenvalues to estimate C.

Error covariance matrix of any linear function y of $\mathbf{v}$ is then obtained as $\mathsf{C}_y = \mathsf{L}_y^T\mathsf{C}\mathsf{L}_y$, where L_y is the matrix representation of the corresponding linear operator $y = \mathsf{L}_y[u,v]$. Of particular importance for applications are the divergence $\mathsf{L}_d[u,v] = \partial_x u + \partial_y v$ and vorticity $\mathsf{L}_c[u,v] = -\partial_y u + \partial_x v$ operators. Mapping the diagonal elements of the corresponding covariance matrices provides an estimate of the accuracy of retrieving these quantities from HFR data, and may give guidance to better configuring the HFR networks.

3. Twin-data experiments

3.1. Setting

To assess the performance of the interpolation scheme, we conducted twin data experiments with simulated HFR observations. Their setting (Fig. 1) was chosen to mimic the real experiment in the Bodega Bay conducted in spring and summer of 2003. The reconstructed velocity fields are sampled along the beams with 2 km discretization, and the azimuthal increment of the beam directions is 5°.

We conducted eight major series of twin-data experiments by varying the interpolation method (2dVar, LI and OMA), the interpolated field (patterns A and B in Fig. 1) and coverage of the domain by observations (with and without the gap shown in the lower left panel in Fig. 1). Within each series of experiments we varied the noise level ν in the simulated observations: radial velocities v_k^* "observed" at points $\mathbf{x}_k$ were defined by adding white noise w to projections of the true currents $\mathbf{v}^t$ on the beam directions $\mathbf{r}_k$:

$$v_k^* = (\hat{P}_k^{ij}\mathbf{v}_{ij}^{\mathbf{t}} \cdot \mathbf{r}_k) + \nu V w \tag{4}$$

Here V is the typical magnitude of $\mathbf{v}^t$ and ν is the scalar parameter whose reciprocal has the meaning of signal/noise ratio. Three values of ν (0.1, 0.3 and 0.5) were tested within each series. In correspondence with Eq. (4), HFR measurement errors in Eq. (3) were defined as $\sigma(v_k^*) = \nu V$.

The "true" current fields $\mathbf{v}^t$ (Fig. 1, upper panels) were defined on a regular 40×32 grid with 2 km step in x and y directions. The fields were obtained by specifying the corresponding stream function ψ and velocity potential φ: first, positions of the jets were set by defining jumps in ψ along the jet axes visible in Fig. 1. In all other spatial locations the stream function was independent of horizontal coordinates. Second, these step-like elongated structures in ψ were smoothed by the biharmonic filter, whose cutoff scale was set to 6 and 10 km for patterns A and B, respectively. Third, Gaussian-shaped eddies of different size and strength were placed between the jets (nine for flow pattern A and three for B). Similar technique was used to define φ with the difference that the amplitude of jets and eddies was 5–6 times smaller and all the eddies had the opposite sign to ψ to simulate convergence within anticyclonic structures and vice versa. The resulting currents have the typical velocity magnitude $V = 0.59$ and 0.11 m/s for cases A and B, respectively.

The spatial scales of the "true" currents in Fig. 1 are somewhat smaller than those typically seen in the Bodega Bay (Kaplan and Lekien, 2007; Section 4 of the present manuscript), especially for the case A, where the domain was populated with more numerous eddies and concentrated jets. This was done to test the algorithm in a somewhat more "violent" regime, which could be encountered in other domains.

To assess the impact of gaps in the spatial coverage by HFR observations, we used two simulated data sets: one with the full coverage by the radial velocity data (shown by gray dots in the lower right panel in Fig. 1), and another without observations in the rectangular domain (lower left panel in Fig. 1). These data sets had $K = 2011$ and 1699 observation points, respectively.

After removing coastal grid points and masking out grid points not covered by observations, the total number of the velocity grid point values was $N = 993$, so that the problem had $2N = 1986$ degrees of freedom.

The quality of interpolation was monitored by several parameters. Misfit between the interpolated field and the data m_d was defined as $m_d = \overline{|\nu_k - v_k^*|}/\overline{|v_k^*|}$, where ν_k is the projection of interpolated velocity on the radar beam at the kth measurement point and overline denotes averaging over the measurements. Velocity error e_v was defined as the mean absolute difference between the true $\mathbf{v}^t$ and the interpolated $\mathbf{v}$ currents normalized by V: $e_v = \langle|\mathbf{v}^t - \mathbf{v}|\rangle/V$, where angular brackets denote averaging over the interpolation grid. Similar expressions are used to assess the interpolation qualities e_d, e_c of the divergence and vorticity fields:

$$e_d = \langle|\mathrm{div}(\mathbf{v}^t - \mathbf{v})|\rangle/\langle|\mathrm{div}\mathbf{v}^t|\rangle; \quad e_c = \langle|\mathrm{curl}(\mathbf{v}^t - \mathbf{v})|\rangle/\langle|\mathrm{curl}\mathbf{v}^t|\rangle$$

Dependence of the bogus data inverse variances W_{ij}^d and W_{ij}^c on x and y was defined to be inversely proportional to the number of observation points n_{ij} within the $2\delta x$ radius from a given grid point $\mathbf{x}_{ij}$:

$$W^{c,d} = \alpha^{c,d} n_{ij}^{-1} \tag{5}$$

In the case when no observations were found in the $2\delta x$ vicinity (e.g., in the center of the gap) n_{ij} was set to 1.

The proportionality coefficients α^c and α^d are equal to the inverse squared magnitudes c^{-2} and d^{-2} of the corresponding fields. The typical magnitude c of $\Delta\mathrm{curl}\mathbf{v}$ is estimated as $c = V/(L^2\delta x)$, where L is the spatial scale of the interpolated velocity field. Both V and L are deduced from statistical analysis of the simulated HFR data: for circulations A and B the spatial scales are estimated as $L_A = 5.9$ km and $L_B = 8.8$ km. We further assume that the typical magnitude of the divergence field may be substantially smaller than c (e.g., for quasigeostrophic currents $d \sim \gamma c$, where γ is the Rossby number). In the presented experiments we set $\gamma = 0.2$ according to the definition of the "true" currents shown in Fig. 1. In practice, γ can be assessed with a reasonable accuracy from physical considerations.

Since the result of interpolation depends only on the relative magnitudes of the terms in the cost function, it is convenient to characterize the algorithm by the non-dimensional parameter $W_*^c = \nu^2 V^2 \alpha^c K/4N\delta x^6 = \nu^2 KL^4/4N\delta x^4$ which has the meaning of the ratio between the vorticity regularization and the data misfit terms. Since we assume that $W^d = \gamma^{-2} W^c$, the value of W_*^c entirely defines the shape of the cost function for a given field (A or B) and data configuration (with or without the gap).

Initial values of $\mathbf{v}$ are set to zero in all the interpolation experiments to ensure uniqueness of the solutions in the sense that unobserved patterns are not present in the interpolated fields (Section 2.3). Experiments were conducted on a 2.66 GHz single-processor PC and it took 0.7 s of CPU to perform $N \sim 1000$ iterations of the minimization procedure.

3.2. Results

3.2.1. Experiments with three radars

Table 1 compares the performance of the proposed 2dVar interpolation algorithm with the local and non-local interpolation techniques.

The LI method was implemented in its simplest form: velocity $\mathbf{v}_{ij}$ at a given regular grid point $\mathbf{x}_{ij}$ was computed by minimizing

$$\hat{J}_{ij} = \sum_{k=1}^{N_r} [\mathbf{v}_{ij} \cdot \mathbf{r}_k - \hat{Q}_{ij}^k v_k^*]^2 \to \min_{\mathbf{v}_{ij}}$$

where N_r is the number of radars illuminating the grid point and the operator $\hat{Q}_{ij}^k$ linearly interpolates the radial velocities v_k^* measured by kth radar in the immediate vicinity of $\mathbf{x}_{ij}$ onto that grid point. Interpolation is performed in polar coordinates with the origin at the radar's location. Note that LI method provides a robust solution with two radars only if their beams are far from parallel, and it would not work at all with one radar or in the regions not covered by data.

The non-local OMA method is formally free from these limitations, because it implicitly employs additional smoothness

Table 1
Dependence of the interpolated field parameters m_d, e_v, e_c, and e_d on the reconstructed velocity field and noise level v in the simulated HFR data.

Field type	v	No gap				Gap				W_*^c
		m_d	e_v	e_c	e_d	m_d	e_v	e_c	e_d	L (km), κ
A	0.10	0.07	0.13	0.19	0.76	0.07	0.15	0.23	0.82	0.2
		0.26	0.31	0.38	1.52	0.28	0.39	0.50	1.70	6×10^{-3}
		0.10	0.19	0.32	1.19	0.15	0.33	0.49	1.35	
	0.30	0.24	0.29	0.37	1.69	0.24	0.29	0.41	1.59	1.8
		0.28	0.32	0.39	1.66	0.29	0.41	0.51	1.78	6×10^{-3}
		0.22	0.30	0.42	2.07	0.23	0.41	0.58	1.79	
	0.50	0.41	0.31	0.37	1.59	0.42	0.43	0.53	1.87	20*
		0.31	0.34	0.42	1.69	0.32	0.44	0.56	1.93	6×10^{-3}
		0.34	0.41	0.57	2.99	0.49	0.50	0.69	2.36	
B	0.10	0.08	0.10	0.18	0.74	0.08	0.13	0.23	0.83	1
		0.23	0.23	0.34	1.32	0.24	0.26	0.35	1.40	7×10^{-4}
		0.10	0.15	0.26	1.36	0.14	0.30	0.44	1.52	
	0.30	0.28	0.19	0.30	1.44	0.25	0.27	0.43	1.87	9
		0.25	0.25	0.36	1.47	0.26	0.30	0.41	1.61	7×10^{-4}
		0.23	0.29	0.49	2.93	0.33	0.41	0.59	2.34	
	0.50	0.28	0.27	0.40	1.57	0.40	0.39	0.53	2.40	25*
		0.25	0.29	0.41	1.87	0.34	0.37	0.56	2.56	7×10^{-3}
		0.34	0.42	0.71	4.40	0.51	0.50	0.75	3.38	

The 2dVar, OMA and LI results are shown, respectively, in the upper, middle and lower lines of the table cells. Note that LI errors were computed over subdomains which did not include near-coastal regions and the gap (see Figs. 2–5). Parameters for the 2dVar and OMA interpolation schemes are shown in the right column. The 2dVar experiments with $v = 0.5$ were made under the divergence-free approximation ($W_*^d = 10^6$).

information and takes into account kinematic constraints imposed on the velocity field by the coastlines. We implemented the OMA algorithm using the code from http://www.pmc.ucsc.edu/~dmk/software/openMA. The code has two tunable parameters: the modal cutoff length scale L and the regularization weighting constant κ, which can be considered as the analog of W_*^c. In all OMA experiments listed in Tables 1 and 2 we minimized e_v by varying both L and κ for each value of v. In the 2dVar experiments the only tunable parameter is W_*^c, which was determined as discussed above.

Inspection of the 2dVar results (upper lines in Table 1) shows that the algorithm is robust in the entire range of noise levels and provides a reasonable fit to the data. The minimum value of $m_d = 0.07$ achieved at 10% noise in Table 1 (case A) corresponds in magnitude to observation errors introduced by the noise w. In that sense the capability of HFR measurements to capture the structure of the flows in Fig. 1 is quite remarkable. With the increase of observation noise the 2dVar algorithm loses precision, but still provides a noise-consistent fit to the data in the entire range of v for both A and B background states (cf. column 2 and upper lines in columns 3 and 7 in Table 1).

Reconstruction of the divergence and vorticity fields is less accurate since horizontal derivatives amplify grid-scale noise. This is especially visible in the behavior of e_d, because the divergence field was set to be five times weaker than vorticity by the design of the experiments. Nevertheless, the 2dVar algorithm captures the signs and positions of the major structures in the divergence field at the "practical" noise levels of 0.1 and 0.3 (cf. the lower right panel in Fig. 1 and the upper right panels in Figs. 2 and 3).

Compared to the divergence, the vorticity field is reconstructed with much higher quality: both the amplitudes and the locations of most of the structures are well reproduced even at the noise level of $v = 0.3$ (cf. low left panel in Fig. 1 and upper left panel in Fig. 3). The gap in data coverage is also handled well: the overall increase in the interpolation errors e_v, e_d and e_c in Table 1 is consistent with the fraction of Ω occupied by the gap, whereas the saddle-like structure of the vorticity field in the gap is reconstructed quantitatively.

The velocity patterns generated by the three methods (left panels in Fig. 2) may look similar, but nevertheless differ considerably in terms of the velocity interpolation error e_v. The difference becomes more evident after taking the divergence of the fields (right panels in Fig. 2): since OMA does not impose any smoothness constraint on div$\mathbf{v}$, the corresponding divergence field, although being three times weaker than vorticity, looks rather chaotic with the formal error $e_d = 1.52$. A similar featureless pattern is given by LI, with the difference that divergence in near-coastal areas cannot be estimated at all due to either single-radar measurements or beam-crossing angle limitation (nearly parallel radar beams).

Introduction of the gap in data coverage enhances the difference between the three methods (cf. columns 4–8 and 5–9 in Table 1 keeping in mind that LI errors are computed outside the gap). A probable reason for the difference between OMA and 2dVar is the emergence of the spurious maxima in both vorticity and divergence fields inside the gap in the OMA case (cf. upper and middle panels in Fig. 3). In the 2dVar formulation, vorticity and divergence fields cannot have maxima inside the data-void regions, because their variation within the gap closely approaches the behavior of a harmonic function.

Comparison of the 2dVar, OMA and LI lines in Table 1 shows that OMA code tends to provide less precise fit to the data at 10% noise level without the gap (lines 2 and 11 in Table 1). This can be partly explained by a relatively low number of degrees of freedom (number of modes) involved. In the OMA formulation the number of modes is proportional to $(D/L)^2$, where D is the horizontal size of the domain and L is the cutoff length scale. For the reported experiments, however, the optimal length scale L was close to 5–6 and 7 km for the A and B fields, respectively (column 11 in Table 1). These length scales correspond to 130–230 eigenmodes, whose amplitudes were varied to fit the data in 2011 (1699 with the gap) points. In an attempt to achieve a better fit m_d at low

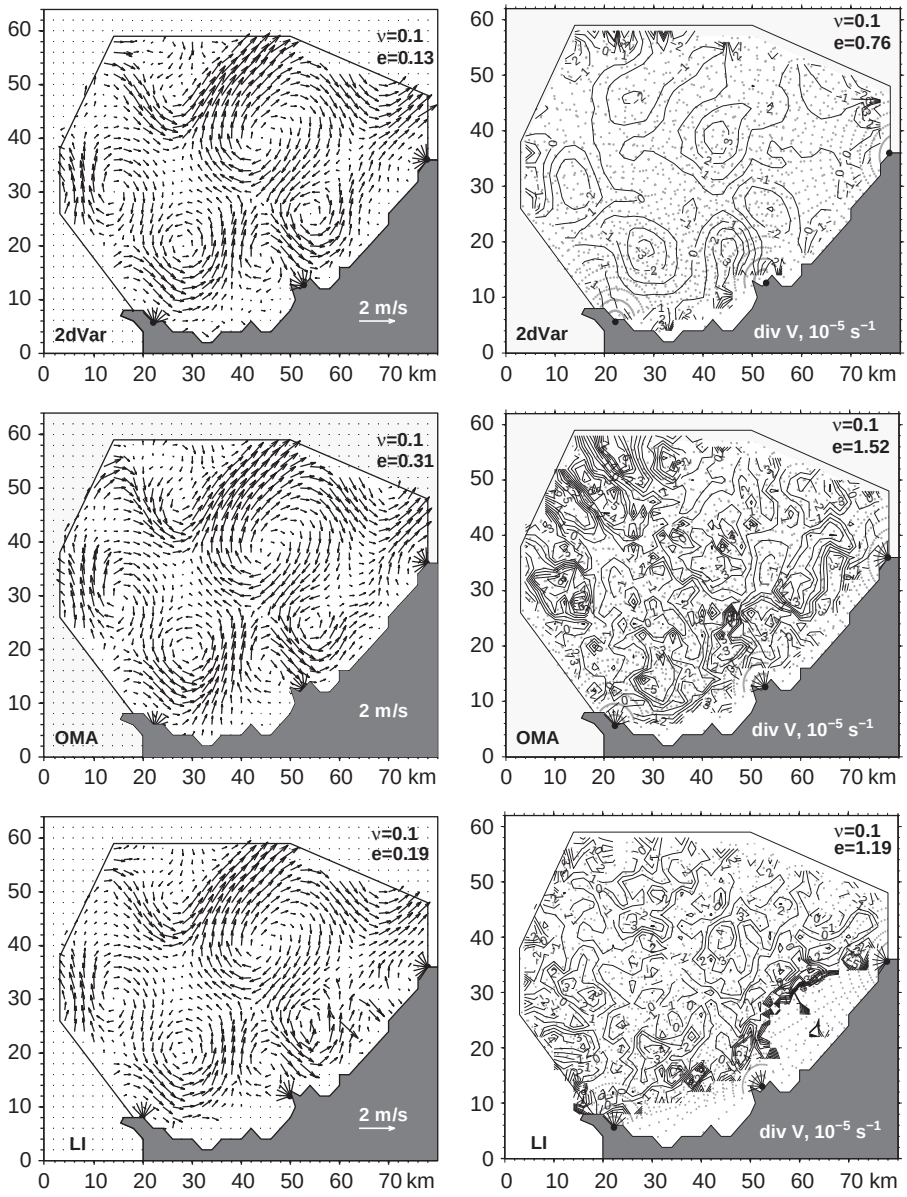

Fig. 2. Comparison between the three methods of interpolation. The noise level v and interpolation errors are shown in the upper right corners of the corresponding panels.

noise levels we tried to increase the number of modes by reducing L to 2–3 km, but that required an increase of regularization parameter κ, otherwise interpolation patterns appeared too noisy, possibly because of the ill-conditioning of the system matrix.

Overall, Table 1 shows that the 2dVar method performs similar to LI and better than OMA at $v = 0.1$ and somewhat better at $v = 0.3-0.5$. When the gap is present in the data describing smaller scale currents (case A, lines 1 and 2, and 4 and 5 in Table 1) 2dVar keeps a significant advantage to OMA up to $v = 0.3$. Both non-local methods (OMA and 2dVar) are better than LI because of their ability to estimate currents within the gap and close to the coastline.

3.2.2. *Experiments with two radars*

In practice, there are often situations when a radar stops operating due to hardware failure or some other reasons. In such case local interpolation methods often fail in a large number of grid points, because they require at least two velocity components for retrieving the velocity vector in a grid cell. The OMA and 2dVar algorithms are essentially non-local and therefore have an ability to interpolate the velocity field over distances exceeding δx.

To investigate the performance of the schemes in such situations, we switched off the rightmost (northern) and/or middle radars and examined the interpolation patterns both with and without the gap in the data. These experiments also allowed us to assess the accuracy of interpolation in the regions where data density was less or close to one observation per grid cell: after removing the northern radar such regions emerge in the upper (western) and right (northern) parts of the domain.

Fig. 4 gives an indication that OMA algorithm is less accurate than 2dVar in such regions: a visual comparison of the vorticity distributions with those in Fig. 1 shows that OMA errors tend to be larger when $y > 40$ km or $x > 60$ km. The LI algorithm performs much worse: derivatives of the velocity field cannot be estimated not only in the region of single-radar coverage ($x > 60$ km) but also near the coastline. Quantitative assessment of e_c for the three

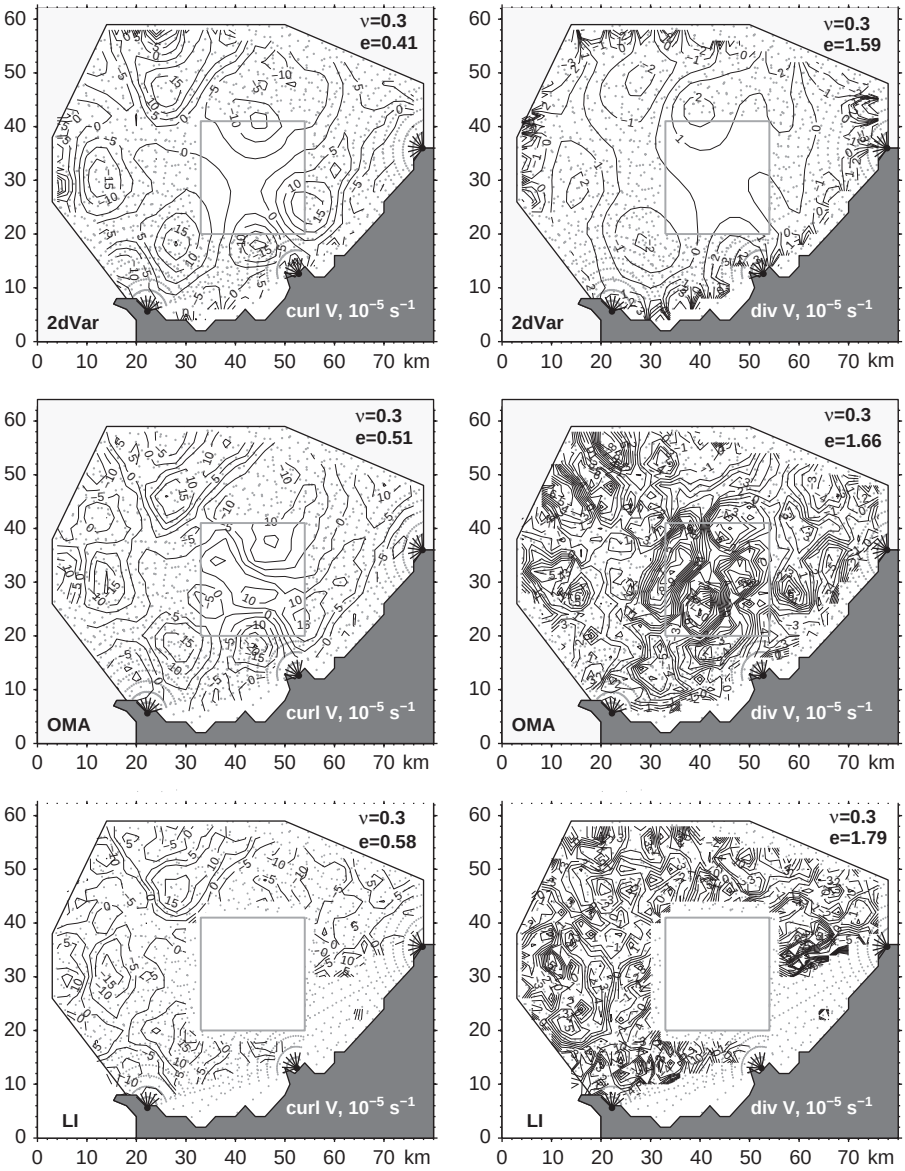

Fig. 3. Comparison between the interpolations of three radars with a gap in observations. The noise level ν and interpolation errors are shown in the upper right corners of the corresponding panels.

interpolation patterns in Fig. 4 also shows that the regions of low data density are the major contributors to the larger value of e_c for the OMA pattern (0.45 vs 0.34 in Fig. 4). The value of e_c for LI (0.39) cannot be objectively compared with these numbers because it was computed by averaging over much smaller area.

Fig. 5 shows an example of interpolation with two radars and the gap in the data. The difference between 2dVar and OMA is already evident from the velocity patterns: OMA produces a spurious jet within the gap which destroys two eddies at the upper and right edges of the data-void region. 2dVar preserves these eddies and the saddle-like structure of the currents within the gap. Comparison of the vorticity fields (right panels in Fig. 5) shows that OMA again produces a maximum inside the gap. The 2dVar pattern appears to be unrealistic inside the gap as well, but compared to OMA has more reasonable structure near the gap's boundary.

The overall results (Table 2) indicate that in the case of two-radar configurations non-local methods (OMA and 2dVar) have a considerable advantage over LI with 2dVar showing somewhat better performance than OMA.

3.2.3. *Experiments with a single radar*

As a matter of curiosity, we also performed experiments with only one radar in operation. The resulting velocity and vorticity fields are shown in Fig. 6. In this experiment we had only 675 observations that were used to retrieve 1986 velocity components in all the grid points. In an attempt to reduce the number of unknowns we reconstructed only the rotational part of the flow by prescribing a large value to $W_*^d = 10^5$ as compared to $W_*^c = 0.1$.

It is interesting that even in this case the interpolation pattern still contains the major circulation features of the "true" field with the formal error level of about 40% (Fig. 6). We attribute this "skill" to specific alignment of the current structures with respect to the radar and to our (correct) prior assumption that the currents are generally non-divergent: the 2dVar algorithm just found a divergence-free field with the smoothest possible

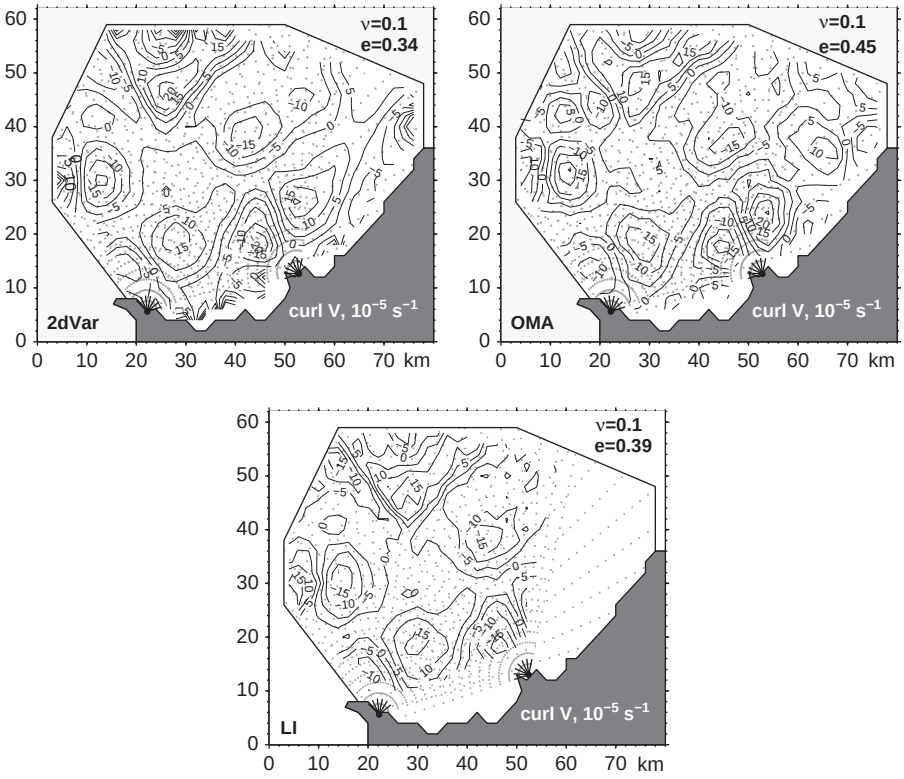

Fig. 4. Comparison between the interpolations with two radars. The noise level v and interpolation errors are shown in the upper right corners of the corresponding panels.

vorticity at the prescribed level of misfit (0.1) with the radial data and zero flow condition at the coast.

3.3. Interpolation errors

Interpolation errors depend on the number of unobserved eigenmodes that are removed from the Hessian spectrum while computing the velocity error covariance matrix C_v. We assumed an eigenmode to be unobserved if the ratio of its eigenvalue to the largest one was less than $\varepsilon = 10^{-7}$, i.e. the relative contribution of the mode to the cost function was comparable with the machine precision. Fig. 7 shows the Hessian spectra computed at different noise levels for observations without (a) and with (b) the gap in the data coverage with three operating radars. As expected, the number of unobserved modes increases with the increase of v and with the reduction of the number of data points in the presence of the gap.

Fig. 8 gives two examples of the interpolation error maps. The left panel shows velocity interpolation errors at $v = 0.1$ in the presence of the gap. To compute the error distribution in Fig. 8a, we took the square roots of the sums of the diagonal elements of C_v that correspond to the error variances for the velocity components u and v, then normalized the map by its maximum value (5.1 cm/s) and took the inverse to obtain the spatial distribution of the relative accuracy of interpolation. The lowest accuracy is attained at the periphery of the domain and near the coastline, where radar beams are almost parallel. A local minimum of the accuracy is also observed inside the gap.

Square roots of the diagonal elements of the divergence error covariance matrix $\mathsf{C}_d = L_d \mathsf{C}_v L_d^{\dagger}$ for the data coverage without the gap are mapped in the right panel. The overall error level

$$\sigma(\mathrm{div}\mathbf{v}) = \sqrt{\frac{\langle \mathsf{C}_d(\mathbf{x}_{ij}, \mathbf{x}_{ij})\rangle}{\langle (\mathrm{div}\mathbf{v} - \langle \mathrm{div}\mathbf{v}\rangle)^2\rangle}}$$

is 0.88, which is consistent with the value of $e_d = 0.76$ (Table 1, line 3, column 6), obtained by the direct comparison of the interpolated pattern with the "true" field.

Error analysis is more expensive than interpolation, because of the necessity to compute the Hessian spectrum, and the pseudo-inverse H^{-1}. In our case these procedures required 296 and 210 CPU s, respectively. For presented method, the eigenfunctions of H (i.e. the modal structure of the velocity error covariance matrix) depend both on the shape of the coastline and on the distribution of the data points over the domain. It is therefore necessary to perform the analysis each time when the data coverage substantially changes.

To estimate errors with n OMA modes, one has first to obtain the $n \times n$ matrix A of the covariances between modal coefficients. This procedure involves two matrix inversions and several multiplications by the $n \times K$ matrix U which projects the modes on radial velocities. The velocity covariance matrix C is then obtained via standard transformation $\mathsf{C} = \mathsf{U}^T \mathsf{A} \mathsf{U}$. We did not have the OMA error analysis code in our disposal, and could not assess its computational cost. A rough estimate indicates, however, that it is of the same order as for the 2dVar algorithm and should also be performed each time when the data coverage changes substantially.

4. Interpolation of the real data

As a final exercise we reconstruct velocity field from real HFR observations in the Bodega Bay conducted on the 30 of July, 2003. The velocity in the upper right panel of Fig. 9 can also be found in Kaplan and Lekien (2007). The interpolation parameters for OMA were $L = 5$ km and $\kappa = 10^{-4}$ as in Kaplan and Lekien (2007). The 2dVar algorithm was used in divergence-free approximation ($W_*^d = 10^6$), because at the estimated noise level of $v = 0.35$, velocity scale of $L = 4.4$ km and sampling discretization of 2 km, it

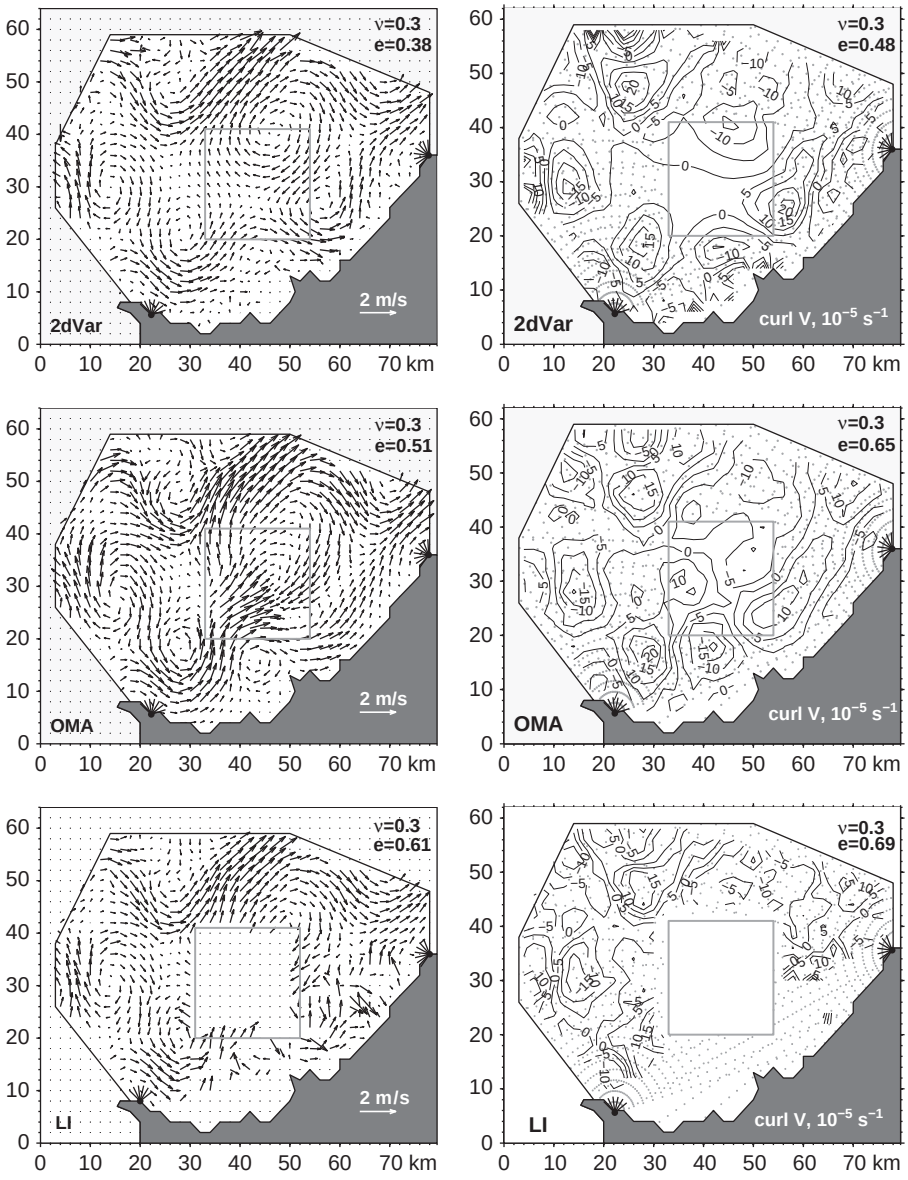

Fig. 5. Comparison between the interpolations of two radars with a gap. The noise level ν and interpolation errors are shown in the upper right corners of the corresponding panels.

Table 2
Same as in Table 1, but for the experiments with the type A circulation and 2-radar configurations.

Radars	ν	No gap				Gap				W_*^c
		m_d	e_ν	e_c	e_d	m_d	e_ν	e_c	e_d	L (km), κ
1.3	0.10	0.09	0.21	0.26	1.31	0.09	0.25	0.35	1.21	0.4
		0.31	0.38	0.48	1.68	0.29	0.48	0.59	2.09	6×10^{-3}
		0.12	0.25	0.42	1.43	0.17	0.39	0.56	1.42	
	0.30	0.24	0.31	0.37	1.60	0.22	0.38	0.48	1.54	1
		0.33	0.42	0.51	2.07	0.31	0.51	0.65	2.37	6×10^{-3}
		0.19	0.39	0.55	2.35	0.24	0.61	0.69	1.82	
1.2	0.10	0.10	0.29	0.34	1.82	0.09	0.32	0.44	1.84	0.6
		0.29	0.36	0.45	1.59	0.30	0.43	0.52	1.88	6×10^{-3}
		0.11	0.32	0.39	1.22	0.15	0.46	0.56	1.27	
	0.30	0.24	0.39	0.45	2.04	0.25	0.44	0.52	2.19	0.8
		0.31	0.41	0.49	2.01	0.30	0.45	0.53	2.03	6×10^{-3}
		0.17	0.42	0.50	2.17	0.25	0.54	0.68	1.26	

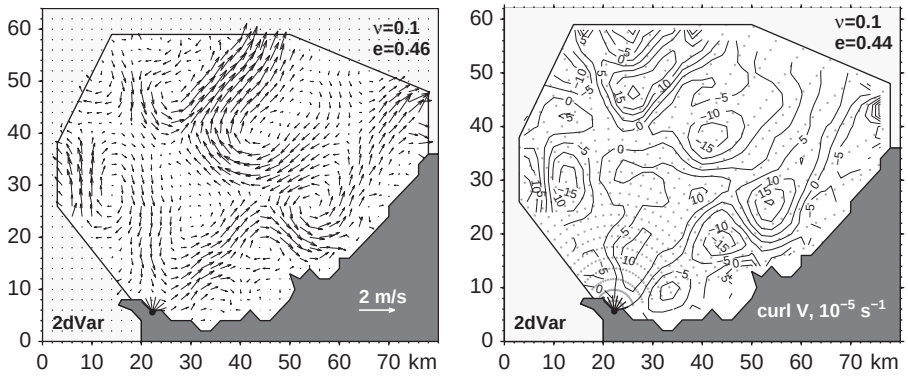

Fig. 6. 2d variational interpolation of the velocity and vorticity fields with only one radar in operation. The noise level ν and interpolation errors are shown in the upper right corners of the corresponding panels.

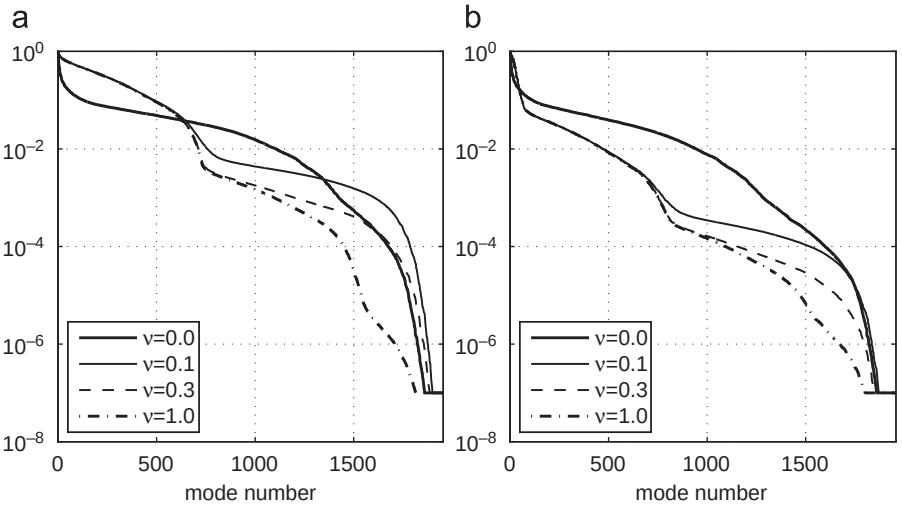

Fig. 7. Normalized Hessian spectra or various ν without the gap in the data (a) and with the gap (b).

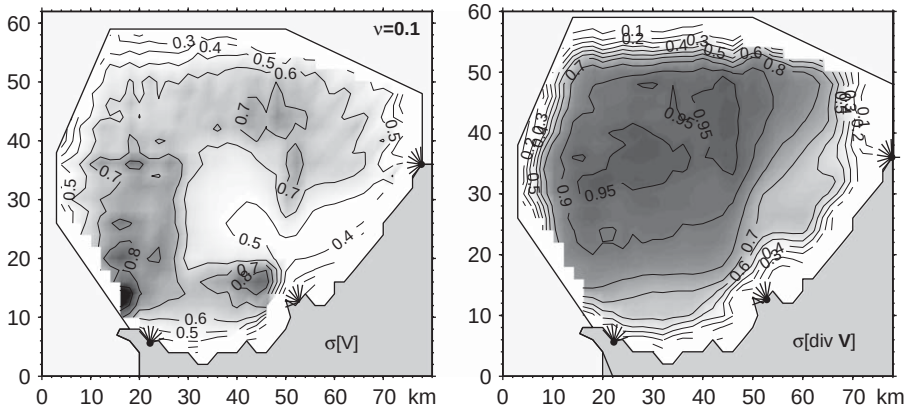

Fig. 8. Accuracy of interpolation for velocity (with the gap in observations, left) and divergence (without the gap, right). Error at a given location can be obtained by dividing 5 cm/s and 2×10^{-5} s^{-1} by the value of the field exposed at the corresponding panel. The observation noise is $\nu = 0.1$.

is hard to obtain statistically confident estimates of the divergence. With 1401 observation points in use $W_*^c = 0.5$.

The LI interpolation (lower panel in Fig. 9) appears to have much lower quality compared to the results of non-local interpolation. This is evident in terms of not only the larger m_d, but also visually: velocity estimates do not exist in the sparsely covered right corner of the domain and are too noisy near the coast where radar beams are close to parallel.

The 2dVar- and OMA-generated patterns are qualitatively similar, although OMA velocity is more smooth and characterized by somewhat larger misfit with the data. The major difference between the 2dVar- and OMA-generated patterns is observed in two regions shown in Fig. 9 by gray rectangles. In the bay between the southern (left) and middle radars the 2dVar pattern shows an indication of anticyclonic circulation, whereas OMA produces a broad offshore current there. The values of e_v computed by

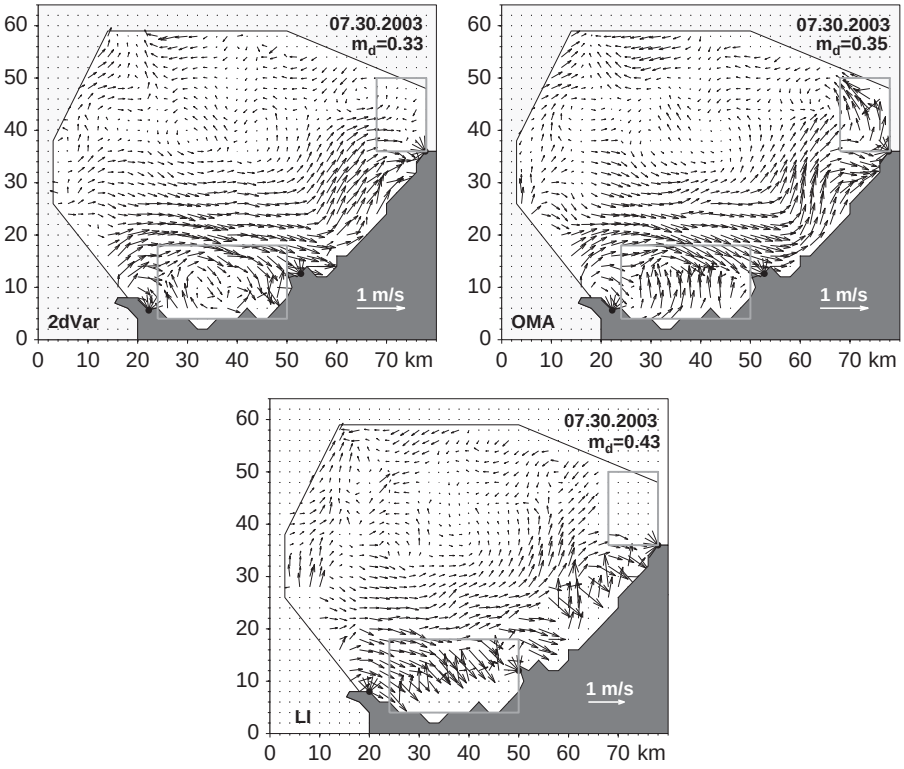

Fig. 9. Velocity field in the Bodega Bay on July 30, 2003, obtained by three interpolation methods. Misfit m_d with the radial velocity data is shown in the upper right corners.

averaging over 265 observation points in that region are 0.26 for 2dVar and 0.27 for OMA, respectively. The region on the right should be considered as a gap, since it contains only four data points near its right boundary.

In the absence of accurate information on the true currents in the region it is hard to estimate the quality of interpolation and compare the results objectively. The reader may evaluate them, however, using a subjective opinion which depends on his/her experience in observing surface currents and/or analysis of their physics.

5. Discussion and conclusions

We have presented a numerical algorithm for retrieving surface velocity field from HFR observations. The algorithm is based on the minimization of a quadratic cost function in the space of all possible configurations of the velocity field. The interpolation problem is regularized by enforcing smoothness in the vorticity and divergence fields. Such approach allows to define the covariance structure of regularization terms using only three scalar parameters that can be easily derived from the statistical analysis of HFR data and basic physical considerations.

In contrast to OMA and the 2dVar algorithms of Park et al. (2007), the proposed method operates in the velocity space, and thus requires less sophisticated operators for the projection of the unknown gridded velocities on the radial components of the current speed. This property provides better conditioning of the system matrix and faster convergence of the minimization algorithm.

Performance of the method is compared with the LI and OMA algorithms. It is shown that the proposed approach is robust and capable to provide a statistically consistent fit to the data in the wide range of signal/noise ratios. The comparison demonstrated similar (to LI) or better (than OMA) performance of the 2dVar technique under relatively high S/N ratios, especially when a 80–90% fit to the velocity field with wide spatial spectrum (containing strong localized features as in type A circulation) is required. At more realistic (less than 3–4) S/N ratios the OMA and 2dVar have similar skill and outperform LI because of their better treatment of the coastal regions where beam directions are close to parallel.

The variational approach gives more flexibility in fitting the data, since the number of modes in use is close to the number of points on the interpolation grid. It is also flexible in the choice of regularization, because the desired smoothness and its spatial variation can be directly controlled by a simple modification of the cost function. The 2dVar method also appears to be more accurate than LI and OMA in the regions with sparse data coverage, and within the large gaps in observations.

Regarding the computational cost, the LI algorithm is significantly better than non-local OMA and 2dVar methods as it required only a fraction of a second of CPU time on a 2.66 GHz PC. Performance of the OMA and 2dVar is compatible: the 2dVar interpolation required 0.7–3 s, whereas the OMA code consumed 2.5–4 s of CPU time on the same machine. The 2dVar error analysis is more expensive (5–8 min of CPU time) than interpolation, since it requires inversion of the Hessian matrix. We did not have the OMA error analysis code in the disposal. However, rough estimates show that the corresponding CPU time should be compatible. In any case, at the present stage of computer technology the computational cost for both non-local methods is sufficiently low to allow their operation in real-time regime.

In formulating the 2dVar algorithm we elected to enforce smoothness in the interpolated patterns of divergence and vorticity, because these quantities are important for physical applications. Because the penalized Laplacians of divergence and vorticity represent high-order derivatives of the velocity field,

stochastic models for $\Delta\mathrm{curl}\mathbf{v}$ and $\Delta\mathrm{div}\mathbf{v}$ are accurate under the diagonal approximation of the corresponding covariance matrices and require definition of only three parameters, which can be inferred from statistical analysis of HFR data. This property of the algorithm allows the user to avoid *ad hoc* definitions of the background velocity field and its covariance that are used to regularize the problem in the method of Park et al. (2007). However, if the background field and its statistics are known with enough confidence (e.g., from a long data-driven model run) the corresponding term could be easily included into the cost function without substantial modifications of the code.

The major idea of the study was to develop an inexpensive algorithm for real-time interpolation of the surface currents at the space/time scales resolved by a typical HFR system. The unresolved processes may induce spatial correlations between the observed radial velocities. These cross-correlations could be taken into the account by specifying the off-diagonal elements of the error covariance matrix σ. The corresponding modification of the code could be done easily as soon as reliable estimates of the spatial error covariances become available. Similarly, the no-slip kinematic constraints can either be relaxed or be replaced by more sophisticated drying/wetting conditions if the HFR-covered area is bounded by gently sloping beaches.

Another straightforward development of the algorithm is refinement of the spatial pattern of the regularization variances W^{-1}: the data density function can be improved by taking into account angles between the beam directions in the vicinity of a grid point. Finally, the interpolation problem can be reformulated in terms of the finite elements that will increase the accuracy in presentation of the coastlines and data-dense regions.

Acknowledgments

This study was supported by the NOAA Grant 09FA1-21 and the Frontier Research system for Global Change through its funding of the International Pacific Research Center. A substantial part of this work was completed during MY visit to the Université du Littoral with support from the French Ministry of Education and Research. We thank David Kaplan for providing us the HFR data in the Bodega Bay and his valuable comments on the OMA code.

References

Breivik, O., Sætra, O., 2001. Real time assimilation of HF radar currents into a coastal ocean model. J. Marine Sys. 28, 161–182.

Chapman, R.D., Shay, L.K., Graber, H.C., Edson, J.B., Karachintsev, A., Trump, C.L., Ross, D.B., 1997. On the accuracy of the HF radar surface measurements: intercomparisons with ship-based sensors. J. Geophys. Res. 102, 18,737–18,748.

Chavanne, C., Janekovic, I., Flament, P., Poulain, P.-M., Kuzmic, M., Gurgel, K.-W., 2007. Tidal currents in the northwestern Adriatic: high-frequency radar observations and numerical model predictions. J. Geophys. Res. 112, C03S21 (doi: 10.1029/2006JC003523).

Eremeev, V.N., Ivanov, L.M., Kirwan, A.D., 1992. Reconstruction of oceanic flow characteristics from quasi-Lagrangian data: 1. Approach and mathematical methods. J. Geophys. Res. 97 (C6), 9733–9742.

Graber, H.C., Thompson, D.R., Carabde, R.E., 1996. Ocean surface features and currents measured with synthetic aperture radar interferometry and HF radar. J. Geophys. Res. 101 (C11), 25,813–25,832.

Gilbert, J., Lemarechal, C., 1989. Some numerical experiments with variable-storage quasi-Newton algorithms. Math. Program. 45, 407–435.

Hisaki, Y., Tokeshi, T., Fujie, W., Sato, K., Fujii, S., 2001. Surface current variability east of Okinawa Island obtained from remotely sensed and in situ observational data. J. Geophys. Res. 106 (C12), 31,057–31,073.

Kaplan, D., Lekien, F., 2007. Spatial interpolation of surface current data based on open-boundary modal analysis. J. Geophys. Res. 112, C12007 (doi: 10.1029/2006JC003984).

Lekien, F., Coulliette, C., 2004. Open-boundary modal analysis: interpolation, extrapolation and filtering. J. Geophys. Res. 109, C12004 (doi: 10.1029/2004JC002323).

Lipphardt, B.L., Kirwan, A.D., Grosch, C.E., Lewis, J.K., Paduan, J.D., 2000. Blending HF radar and model velocities in Monterey Bay through normal mode analysis. J. Geophys. Res. 105 (C2), 3425–3450.

Park, K., Li, Z., Farrara, J., Chao, Y., 2007. Assimilation of high-frequency radar measurements for coastal currents using ROMS. Proc. SPIE 6685, 55–68.

Prandle, D., 1993. Year-long measurements of flow through the Dover Strait by HF radar. Oceanol. Acta 16, 457–468.

Sentchev, A., Yaremchuk, M., 2007. VHF radar observations of surface currents off the northern Opal coast in the eastern English Channel. Cont. Shelf Res. 27, 2449–2464.

Shay, L.K., Martinez-Pedala, J., Cook, T.M., Haus, B.K., 2007. High-frequency radar mapping of surface currents using WERA. J. Atm. Oceanic Tech. 24, 484–503.

Thacker, W.C., 1988. Fitting models to inadequate data by enforcing spatial and temporal smoothness. J. Geophys. Res. 93, 10,556–10,566.

Thacker, W.C., 1989. The role of the Hessian matrix in fitting models to measurements. J. Geophys. Res. 94 (C5), 6177–6196.