Physica D 351–352 (2017) 42–52

Contents lists available at ScienceDirect

Physica D

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

Optimal strategies for the control of autonomous vehicles in data assimilation

D. McDougall [a,*], R.O. Moore [b]

[a] 201 E. 24th St., Stop C0200, Institute for Computational Engineering and Sciences, University of Texas at Austin, Austin, TX 78712, United States

[b] Department of Mathematical Sciences, New Jersey Institute of Technology, University Heights, Newark, NJ 07102, United States

HIGHLIGHTS

- Propose a sequential method for optimal control.
- Method is coupled with the Kalman filter.
- Show decreased posterior uncertainty with optimally controlled gliders.

ARTICLE INFO

Article history:
Received 7 December 2015
Received in revised form
28 February 2017
Accepted 2 April 2017
Available online 5 May 2017
Communicated by H.A. Dijkstra

Keywords:
Bayesian inverse problem
Lagrangian data assimilation
Optimal control
Ocean glider

ABSTRACT

We propose a method to compute optimal control paths for autonomous vehicles deployed for the purpose of inferring a velocity field. In addition to being advected by the flow, the vehicles are able to effect a fixed relative speed with arbitrary control over direction. It is this direction that is used as the basis for the locally optimal control algorithm presented here, with objective formed from the variance trace of the expected posterior distribution. We present results for linear flows near hyperbolic fixed points.

1. Introduction

The need for a more accurate and better resolved estimate of oceanic flows is being driven by a number of pressing global issues, including the crisis facing many species of fish and waterborne organisms, the mitigation of pollutants resulting from spills and offshore contamination, and the important role played by ocean dynamics on climate change. Scientific efforts to estimate ocean flow began in the 1980s with the work of Robinson [1], but have enjoyed limited success due to a lack of observational data. In an effort to improve the current state of understanding of the world's oceans, autonomous vehicles (AVs) are being deployed for the collection of physical oceanography data in a growing number of projects around the globe. One example of AVs is autonomous underwater vehicles (AUVs), which are equipped with adjustable wings that convert vertical momentum induced by battery-powered changes in buoyancy into forward momentum.

Notwithstanding a myriad of challenges caused by their immersion in a high-inertia environment and the limited power and bandwidth available for communication, the effectiveness of AVs comes from their ability to remain in the field for prolonged deployment and from their capacity for controlled self-propulsion. To understand why this control is so important, one need simply consult the growing literature on filtering of complex systems. The dimension associated with a meaningful estimate of an oceanic velocity field is far larger than a typical AV cohort size, so that the observations are necessarily sparse in space. They may also be sparse in time for autonomous underwater vehicles (AUVs) that are unable to communicate while submerged. Even large cohorts of Lagrangian drifters that simply advect with the flow are often subject to collapse onto a small subset of stable dynamic features of the flow, rendering their observations increasingly less informative with time. Recent demonstrations that such drifters face inherent information barriers suggest that simply increasing the size of a

* Corresponding author.
E-mail addresses: damon@ices.utexas.edu (D. McDougall), rmoore@njit.edu (R.O. Moore).

http://dx.doi.org/10.1016/j.physd.2017.04.001

deployed host is not effective without implementing a method to allow them to cross flow barriers [2].

To better understand this issue, we reconstruct a two-dimensional flow field from observations taken by AVs that are capable of limited locomotion in addition to advecting with the flow. We refer to these AVs as *gliders*, in distinction to passive drifters. They are capable of measuring the surrounding fluid velocity either directly using Doppler velocimetry or indirectly using sequential position measurements, depending on the instrumentation on board. In this work, we focus on the former case and use the velocity observations to perform a numerical experiment based on a perfect model, with the goal of constructing an estimate of the true parameterized flow field function from observations, along with an associated uncertainty. Our approach to control follows the underlying philosophy of [3] that mesoscale ocean flow fields are dominated by coherent features, such as jets and eddies, which not only dictate the sites where the most informative observations can be taken to estimate these structures, but also provide the most effective transport mechanisms for navigation of the physical domain.

If the ocean drifter can be controlled to move into and through these structures then the information gained should be richer in terms of capturing the key properties of the flow field. Whereas in [3] this was shown by utilizing an ad hoc control strategy, the present work attempts to put this philosophy on a more systematic footing through the application of an *optimal* control strategy. Each optimal control calculation is inserted between successive observations in a sequential filter, utilizing the posterior field estimate to compute a control that most effectively minimizes the next expected posterior variance.

In general, reconstructing the flow requires one to solve an inverse problem which is most naturally posed in a Bayesian framework [4], the solution to which is a probability distribution on the appropriate function space or parameter space, depending on the velocity field's representation. Kalman filters and Kalman smoothers [5–10] completely quantify this distribution by its first two moments in the case of linear processes and Gaussian initial distributions, and are popular approximate methods for nonlinear flows and non-Gaussian distributions. The Kalman filter is the assimilation mechanism used here, ostensibly due to the use of linear stationary velocity fields and Gaussian distributions but also to provide the straightforward extension to weakly nonlinear flow models through the extended Kalman filter [4].

Alternative approaches include variational methods that view the solution not as a distribution but as the argument of a cost function that is optimized [11–19], and particle filtering methods [20,21] that approximate the continuous posterior random variable by a discrete set of state realizations (particles) with associated weights. Updating the particles and weights as new observations are made is difficult and can lead to degenerate posteriors in high-dimensional problems [22]. Sampling methods utilize the Metropolis–Hastings algorithm [23,24] to sequentially generate correlated samples from the posterior distribution [25–40]. Sampling methods are ideal for the online computation of moments through the use of unbiased estimators, but are not suited to applications where the data arrive sequentially, and are expensive in settings where the distribution must be used in an inverse problem to determine the minimizing control.

The approach of 'adaptive observations' or 'targeted observations' is very similar. The basic idea is that one may be able to engineer where to make an observation to optimize some objective. Berliner, Lu, and Snyder [41] formulate the statistical design problem of adaptively choosing observations for the purposes of improving a prediction. Furthermore, it establishes the importance of the choice of objective to optimize. The work of Hamill and Snyder [42] choose the trace of the posterior covariance matrix in an Ensemble Kalman Filter, and operate in the context of a finite network of observations from which a single observation location is chosen. An adjoint sensitivity approach to targeting observations is proposed in Daescu and Navon [43] for the variational 4D-Var method. They target observations in time as well as in space to optimize the model forecast error and do not claim to find optimal observation paths. The Ensemble Transform Kalman Filter [44], later improved upon in [45], is a filter that also offers an efficient method of estimating how a particular observation affects the forecast error covariance mathematically. In all of the targeted observation literature hitherto mentioned, assessing the efficacy of each targeting strategy is done within the context of an observation system simulation experiment (OSSE). The work presented here differs from previous targeted observation work in two ways. First, the observations cannot be arbitrarily targeted and must arise from positional data of ocean gliders. Second, the formulation presented is in the context of optimal control and as such can be done 'online' within a Kalman Filter resulting in (locally) optimal observation paths.

Section 2 describes the setup of the problem under consideration in a relatively broad setting, including the basic flow assumptions and assimilation model, followed by a derivation of the method used to find local minimizers of the objective. Section 3 identifies an existence problem with the variational formulation and introduces a simple approach to address it. Section 4 provides numerical results for the inference of four linear time-independent velocity fields modeling flow near hyperbolic fixed points with different stability properties. Section 5 offers conclusions and discusses extensions to the methodology. The Appendix describes an alternative method for obtaining locally optimal gliders paths used for comparison purposes and a convergence study.

2. Setup

The context of this study is oceanographic data assimilation, where the goal is to estimate a velocity field conditioned on sparse observational data recorded by Lagrangian *gliders* whose positions evolve according to the flow being assimilated in addition to a modest capacity for self-propulsion. The gliders therefore move according to

$$\dot{\boldsymbol{z}}^{(k)} = \boldsymbol{v}(\boldsymbol{z}^{(k)}(t), t) + \boldsymbol{u}^{(k)}(t), \quad k = 1, \dots, K, \tag{1}$$

with $\boldsymbol{z}^{(k)} \in \mathbb{R}^2$ containing the position of glider k in a 2-dimensional domain and $\boldsymbol{v}(\boldsymbol{z}^{(k)}, t) := (v_1(\boldsymbol{z}^{(k)}, t), v_2(\boldsymbol{z}^{(k)}, t))^\top$. The flow evolves according to a model expressed as

$$\frac{\partial \boldsymbol{v}}{\partial t} = \boldsymbol{f}(\boldsymbol{v}) + \boldsymbol{\xi}(\boldsymbol{x}, t), \tag{2}$$

where $\boldsymbol{\xi}$ is a term representing process noise given by

$$\mathbb{E}(\boldsymbol{\xi}(\boldsymbol{x}, t)\boldsymbol{\xi}^\top(\boldsymbol{x}', t')) = \boldsymbol{Q}(\boldsymbol{x} - \boldsymbol{x}')\delta(t - t'). \tag{3}$$

The kth glider's self-propulsion is expressed by control $\boldsymbol{u}^{(k)}(t)$, where $\boldsymbol{u}^{(k)}(t) \equiv 0$ represents the uncontrolled case (i.e., Lagrangian drifters [46]) and $|\boldsymbol{u}^{(k)}(t)| \equiv u_{\max}$ represents the case considered here with the relative speed of gliders fixed at $u_{\max}$ [47].

The direct velocity observations considered here take the form

$$\boldsymbol{d}_i^{(k)} = \boldsymbol{v}(\boldsymbol{z}^{(k)}(t_i), t_i) + \eta_i^{(k)}, \quad k = 1, \dots, K, \tag{4}$$

where $\eta_i^{(k)} \sim \mathcal{N}(\boldsymbol{0}, \boldsymbol{R}_i)$ is i.i.d. Gaussian-distributed measurement noise and $\boldsymbol{d}_i^{(k)} \in \mathbb{R}^2$. The mapping between "state space" (estimates of $\boldsymbol{v}(\boldsymbol{x}, t)$) and "observation space" (values of $\boldsymbol{d}$) is referred to as the observation operator which, in this case, is formally given by the convolution

$$\boldsymbol{H}(\boldsymbol{v}) = \begin{pmatrix} \delta(\boldsymbol{x} - \boldsymbol{z}^{(1)}(t_i)) * \boldsymbol{v} \\ \delta(\boldsymbol{x} - \boldsymbol{z}^{(2)}(t_i)) * \boldsymbol{v} \\ \vdots \\ \delta(\boldsymbol{x} - \boldsymbol{z}^{(K)}(t_i)) * \boldsymbol{v} \end{pmatrix}. \tag{5}$$