Small and Adrift with Self-Control: Using the Environment to Improve Autonomy

M. Ani Hsieh, Hadi Hajieghrary, Dhanushka Kularatne, Christoffer R. Heckman, Eric Forgoston, Ira B. Schwartz, and Philip A. Yecko

Abstract We present information theoretic search strategies for single and multi-robot teams to localize the source of a chemical spill in turbulent flows. In this work, robots rely on sporadic and intermittent sensor readings to synthesize information maximizing exploration strategies. Using the spatial distribution of the sensor readings, robots construct a belief distribution for the source location. Motion strategies are designed to maximize the change in entropy of this belief distribution. In addition, we show how a geophysical description of the environmental dynamics can improve existing motion control strategies. This is especially true when process and vehicle dynamics are intricately coupled with the environmental dynamics. We conclude with a summary of current efforts in robotic tracking of coherent structures in geophysical flows. Since coherent structures enables the prediction and estimation of the environmental dynamics, we discuss how this geophysical perspective can result in improved control strategies for autonomous systems.

M. Ani Hsieh, Hadi Hajieghrary, and Dhanushka Kularatne
Scalable Autonomous Systems Lab, Drexel University, Philadelphia, PA 19104, USA e-mail: {mhsieh1,hadi.hajieghrary,dnk32}@drexel.edu

Christoffer R. Heckman
Autonomous Robotics & Perception Group at the University of Colorado, Boulder, CO 80309 USA e-mail: christoffer.heckman@colorado.edu

Eric Forgoston
Department of Mathematical Sciences, Montclair State University, Montclair, New Jersey 07043, USA e-mail: eric.forgoston@montclair.edu

Ira B. Schwartz
Nonlinear Systems Dynamics Section, Plasma Physics Division, Code 6792, U.S. Naval Research Laboratory, Washington, DC 20375, USA e-mail: ira.schwartz@nrl.navy.mil

Philip A. Yecko
Albert Nerken School of Engineering, The Cooper Union, New York, NY 10003, USA e-mail: yecko@cooper.edu

1 Introduction

The motions of small vehicles are more significantly impacted by the environments they operate in, allowing the environment to be leveraged to improve vehicle control and autonomy, particularly for aerial and marine vehicles [16]. Consider the example of a small autonomous marine vehicle (AMV) operating in the turbulent ocean. The tightly coupled vehicle and environmental dynamics makes control challenging, but offers nearly limitless environmental forces to be exploited to extend the power budgets of small, resource constrained vehicles.

Recent work in the marine robotics community has verified that AMV motion planning and adaptive sampling strategies can be improved when accounting for dynamics of the fluidic environment [13, 28, 27, 16, 17, 10]. Further progress is hindered by the immense complexity of the atmospheric and/or the ocean dynamics, which involves the interplay of rotation, stratification, complex topography and variable thermal and atmospheric/oceanic forcing, not to mention thousands of biological, chemical, and physical inputs. Theoretical and experimental efforts to model atmospheric and ocean flows have made progress with the help of simpler, so-called "reduced" models, but these models are too idealized and limited in applicability for use in the field. On the other hand, atmosphere and ocean hindcasts, nowcasts, and forecasts provided by National Oceanic and Atmospheric Administration (NOAA) and Naval Coastal Ocean Model (NCOM) [22] and regional ocean model systems (ROMS) [28] include data assimilated from satellite and field observations. The overall quality of these data is highly dependent on how well a given region of interest is instrumented [25, 26], stymieing attempts to incorporate historical and/or forecasted flows into vehicle motion planning and control strategies.

Fortunately, while geophysical flows are physically complex and naturally stochastic, they also exhibit coherent structure. Fig. 1a and 1b[1] show examples of coherent structures that are easily discerned from the snapshots of atmospheric and ocean surface currents. Knowledge of coherent structures enables prediction of flow properties, including transport, where they are known to play a key role. The Gulf Stream is a prominent example of a mesoscale coherent jet whose heat transport is a critical component of global weather and climate. In addition to large mesoscale eddies and jets, smaller, sub-mesoscale, coherent features such as fronts, often grow from instabilities of the larger scale motions [15]. These features impact both transport and the ocean's *primary production*, and thus storage, of organic matter. Using geophysical fluid models, details from flows such as the Gulf Stream can be used to diagnose the underlying geophysical fluid dynamics. This, in turn, enables the prediction of various physical, chemical, and biological processes in general geophysical flows. More importantly, predictions based on coherent structures may be exploited more effectively than the detailed predictions offered by state-of-the-art numerical models. Coherent structures can thus provide a basis from which one can construct a vastly reduced order description of the fluid environment.

[1] For full animation visit `http://earth.nullschool.net/` and `http://svs.gsfc.nasa.gov/vis/a000000/a003800/a003827/`.

The challenge of robotic control in a fluidic medium is closely tied to the dual problems of mixing and optimal sensing. G.I. Taylor described and quantified how even simple steady shear flows enhance the mixing of a contaminant in a fluid, a process now known as *Taylor dispersion* [12]. In its simplest form, Taylor dispersion is due to the creation of small scales which locally enhance contaminant gradients and thus optimize diffusion. While achieving better mixing over long times, for shorter time scales, the ability of flows to form structures may also inhibit mixing, as when contaminant is trapped in a coherent vortex.

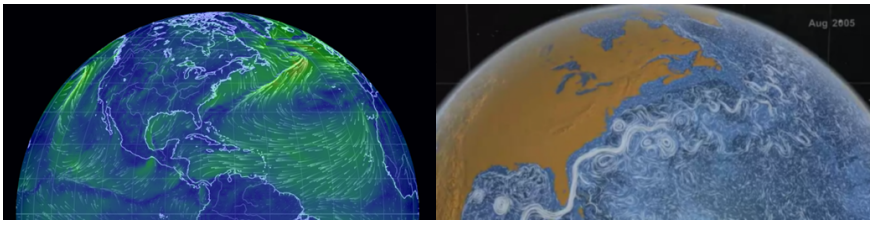

(a) (b)

Fig. 1: (a) Visualization of atmospheric currents for January 2015 using data provided by Global Forecasting System, NCEP, National Weather Service, and NOAA. (b) Visualization of ocean surface currents for June 2005 through December 2007 using NASA/JPLs Estimating the Circulation and Climate of the Ocean, Phase II (ECCO2) ocean model.

In this work we examine autonomous sampling of a dispersive event in a geophysical flow. In particular, we consider the problem of finding the source of a contaminant leak in a turbulent flow and discuss how search strategies can be significantly improved from a GFD perspective. The GFD framework is based on a specific class of coherent structures, called Lagrangian coherent structures (LCS). LCS are similar to separatrices that divide a flow into dynamically distinct regions; hyperbolic LCS can be understood as extensions of stable and unstable manifolds to general time-dependent flows [6, 5, 4].

In our ongoing work we have tracked LCS using teams of autonomous robots in geophysical fluidic environments, while Inanc et al. showed that time and fuel optimal paths in the ocean can coincide with LCS boundaries [13, 23]. As such, knowledge derived from geophysical fluid dynamics (GFD) can significantly improve the overall quality of vehicle autonomy. For example, planning energy efficient trajectories, maintaining sensors in their desired monitoring regions [17, 10, 7], and enabling computationally tractable and efficient estimation and prediction of surrounding environmental dynamics. We claim that LCS based flow knowledge can be used to improve contaminant tracking strategies and support our claims using geophysical fluid and dispersion models. We conclude with a discussion of major challenges and opportunities ahead.

2 Contaminant Source Localization in Turbulent Mediums

To illustrate these ideas, consider the problem of finding the source of a contaminant or hazardous waste plume in a turbulent medium. Turbulence poses significant challenges for the localization and tracking of material dispersion sources since it breaks up continuous patches of material into seemingly random moving disconnected components. As such, gradient-based search strategies based on chemical concentrations become highly unreliable in turbulent mediums since the mixing dynamics renders any estimation of chemical gradients ineffective [21].

In this section, we formulate the source seeking/plume source localization problem as an information theoretic search strategy. The proposed strategy builds upon [30] where the search strategy consists of making moves that maximize the change in entropy of the posterior distribution of the source location. Similar to [8, 1], we rely on a particle filter representation of the posterior belief distribution to make the strategy more computationally viable in large complex spaces and distributable for mobile sensing teams. The main contribution is extending existing state-of-the-art information theoretic search strategies for robots operating in a turbulent flows.

Background and Assumptions We assume the mean rate of chemical detection at position r resulting from a leakage at r_0 follows a Poisson distribution given by:

$$\mathscr{R}(r|r_0) = \frac{R}{\ln(\frac{\lambda}{a})} e^{\frac{(y_0-y)V}{2D}} K_0\left(\frac{|r-r_0|}{\lambda}\right), \tag{1}$$

where $\lambda = \sqrt{D\tau/(1+V^2\tau/4D)}$, R is the emission rate of the source, τ is the finite lifetime of a chemical patch before its concentration falls out of the detectable range, D is the isotropic effective diffusivity of the medium, V is the mean velocity of the current, and $K_0(\cdot)$ is the modified Bessel function of the second kind [30]. The probability of registering the presence of a chemical patch by a sensor depends on the distance and the angle of the sensor from the source, *i.e.*, $(r-r_0)$.

In the search for the source, the set of chemical detection encounters along the search trajectory carries the only available information about the relative location of the source with respect to the robot. Using this information, a robot can construct the probability distribution of the source location using Bayes' rule $P(r_0|\mathscr{T}_t) = P(\mathscr{T}_t|r_0)P(r_0)/\int P(\mathscr{T}_t|r)P(r)dr$. Here, $\mathscr{T}_t$ encapsulates the history of the uncorrelated material encounters along the robot's search trajectory, with $P(\mathscr{T}_t|r_0)$ denoting the probability of experiencing such a history if the source of the dispersion is at r_0. We assume the probability of detecting a material or chemical plume at each step is independent and use Poisson's law to estimate the probability of detecting such a history of the material presence along the search trajectory as:

$$P(\mathscr{T}_t|r_0) = \exp\left(-\int_0^t \mathscr{R}(r(\bar{t})|r_0)\,d\bar{t}\right) \prod_i \mathscr{R}(r(t_i)|r_0), \tag{2}$$

where $r(t)$ is the search trajectory, and $r(t_i)$ is the position of each detection along the trajectory [30]. We note that the assumption of the independence of detections

holds since the location of the source is unknown. Furthermore, chemical plumes quickly disintegrate into disconnected patches in turbulent mediums whose dispersion dynamics are highly nonlinear and stochastic. Rather than attempt to model the complex process dynamics, we assume detection events are independent.

We note that (1) provides an observation model and there is no process model beyond what we assume of the environmental dynamics and the propagation of the source in the medium. As such, the current estimate of the belief distribution of the source location, $P_t(r)$, is used to determine the expected number of positive sensor measurements at the new position, *i.e.,* $\mathscr{R}(r(t)|r_0)$. We assume robots move on a grid within the workspace and at each time step can travel an upper bound of n moves. The maximum number of moves on the grid is determined based on the limitations and bounds on the control input and is based on the vehicle's dynamics. The workspace is assumed to be bounded but obstacle-free.

The objective is to localize the source of the dispersion in the workspace. Since the rate of chemical encounter is dependent on the robot's position with respect to the source, the proposed strategy should result in robots maximizing the expected rate of information acquisition on the source's position.

The Single Robot Search Strategy The information gathered through the chemical encounters shapes the probability distribution function that describes the possible source locations, denoted by $P_t(r)$. Thus, it makes sense to consider a search strategy that drives the robot in a direction that promises the steepest decrease in the entropy of this distribution. The expected rate of information gain at each search step is the expected change in the entropy of the estimated field is given by:

$$\mathbf{E}[\Delta S_t(r \mapsto r_j)] = P_t(r_j)[-S_t] + [1 - P_t(r_j)][(1-\rho(r_j))\Delta S_0 + \rho(r_j)\Delta S_1], \quad (3)$$

where, $S_t = \int P_t(r)\log(P_t(r))dr$ is the Shannon entropy of the estimated field. The first term of (3) corresponds to the change in entropy upon finding the source at the very next step. If the robot successfully localizes the source at the next step, then the change in entropy would be zero and thus end the search.

Using the current belief distribution as the best estimate of the source location, the expected number of positive sensor hits at location r is given by:

$$h(r) = \int P_t(r_0)\mathscr{R}(r|r_0)dr_0 \quad (4)$$

and the probability of a single positive detection follows the Poisson law $\rho(r) = h(r)\exp(-h(r))$. Thus, the second term of (3) accounts for the case when the source is not found at r_j and computes the expected value of the information gain by the robot moving to this new position. At each new location, we assume the robot takes one measurement with its sensor resulting in two possibilities – the robot will have either a positive sensor hit or not. To find the expected value of the change in the utility function, *i.e.,* the entropy, (3) calculates the change in entropy for each possible move and each possible sensor outcome after the move.

The belief distribution for the source location $P_t(r)$ is maintained over all possible source positions. At each step, the robot chooses to move in a direction that enables it to acquire more information and decreases the uncertainty of the source position estimate. Storing and representing the belief distribution becomes computationally challenging especially when the search space spans large physical scales and/or contains complex geometry. This is especially true if robots rely on a fine grid map to calculate the log-likelihood of the belief distribution. Furthermore, since the search hypotheses are often spread across the workspace, closed form descriptions for the belief distribution may not be accessible, especially in geometrically complex spaces. Thus, we employ a particle filter approach to the representation of the belief distribution with limited numbers of randomly drawn particles.

In this work, we assume each robot stores the estimated belief distribution of the source location, $P_t(r)$, using a manageable number of particles, $\{\hat{r}_i, \omega_i\}$, with $\hat{r}_i$ representing the hypothesis for the state (position) and ω_i representing the corresponding weight (probability) of hypothesis i. The probability mass function represented by these set of particles is mathematically equivalent to the sum of the weighted spatial impulse functions [2]:

$$\hat{P}_r(t) \approx \sum_i \omega_i \delta(\hat{r}_i - r), \tag{5}$$

where $\hat{r}_i$ is a hypothesis that survived the re-sampling procedure of the previous step, and the weights, ω_i, of the particles are modified as follows:

$$\omega_i(t) = \omega_i(t-1) e^{-\left(R(r(t)|r_0)\right)} \left(R(r(t)|r_0)\right)^{hit} \tag{6}$$

with $hit = 1$ in the case of a positive detection and $hit = 0$ otherwise. To calculate the entropy of the particle representation of the belief distribution, we use the approach presented in [8] where $S \approx -\sum_{k=1}^{N} w_{(t-1),k}^{(i)} \log\left(w_{(t-1),k}^{(i)}\right)$.

A robot's control decision is determined by the expected change in entropy of the source position estimate. At every time step, the expected information gain from an observation at the next probable robot position is calculated. Since robots are constrained to a maximum of n moves on the grid, they can quantify the expected information gain on the source's position as it moves. If no particles are within the robot's set of reachable points on the grid, then any position that is n moves away from the robot's current position can be chosen as the next step.

Two approaches are used to represent the information gathered during the search. The first one assumes the robot has a limited field of view with no knowledge of the search area. The particles are placed in the robot's coordinate frame, and the control decision is to move in the direction the robot expects to acquire more information. Although the source location may not initially be within the robot's field of view, we assume its sensing range is large enough to provide a good enough measurement for it to determine a direction to move in. The second approach assumes the robot knows the boundary of the search area and can localize itself within the space. Under

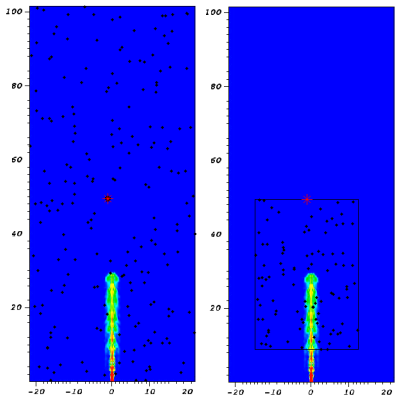

(a) (b)

Fig. 2: The particles representing the hypotheses for the source position. The background shows the gas volume concentration in the water due to an oil spill. Blue denotes low concentration and red denotes high concentration. Details of the plume simulation can be found in [3].

this scenario, the particles used to estimate the source positions are spread over the entire workspace. Fig. 2 depicts the two representations.

It is important to note that we assume robots can measure the flow direction at its current position. This is important for the robot to discern the direction the material plumes are coming from. We note that the weight update given by (6) is different from most particle filter implementations. This is because when localizing the source of a dispersion, re-sampling serves the role of integrating past information into the current estimate of the source position. Therefore, one must take the likelihood of the detection history, (2), into account during the update and eliminate the less probable hypotheses when re-sampling. This is the fundamental difference between the proposed strategy and existing information theoretic strategies for aerial and ground vehicles operating in static environments. The single robot search strategy is summarized in Algorithm 1.

The Multi-Robot Search Strategy To speed up the search strategy, we extend the proposed single robot search strategy to a robot team. This effectively increases the chances of positive encounters and decreases the expected time needed to localize the source. To achieve this, we assume robots can communicate with one another and thus the team can build a shared belief distribution of the source position. While communications can be severely bandwidth limited in a fluidic medium, e.g., underwater environments, the amount of data that must be communicated is low. Robots are only required to exchange their position information when they have detected a material plume at their current location. Since robots initialize with the same belief distribution, each robot moves in a direction that reduces the entropy of this common belief distribution. Such a coordination strategy ensures every individual in the team searches for a single source location. To take into account the impact of robot

input : CURRENT estimate of source position belief distribution
output: UPDATED estimate of source position belief distribution

for All particles within the robot's reachable set **do**
    Update particle weights if the robot moves to the position of the particle, and;
    (1) The robot detects a plume in the new position;
    (2) The robot does not detect a plume in the new position;
    Calculate entropy of the particle filter for (1) and (2);
    Calculate expected reduction in entropy if robot moves to each particle position;
end
Move to the location of the particle with steepest expected reduction in entropy;
Obtain sensor reading and compute new particle weights for the samples;
Re-sample;

Algorithm 1: Single robot search strategy.

motion uncertainties on particle positions, we represent each particle position with a Gaussian distribution and thus employ a Gaussian Mixture Model to represent the probability distribution of the source given by (5). As such, we employ the steps described in [11, 8, 1] to calculate the information utility function:

$$S \approx -\int_{\theta} \Big\{ \sum_{k=1}^{N} w_{(t-1),k}^{(i)} p(\theta_t | \theta_{t-1} = \hat{\theta}_{(t-1),k}^{(i)}) \log \big(\sum_{k=1}^{N} w_{(t-1),k}^{(i)} p(\theta_t | \theta_{t-1} = \hat{\theta}_{(t-1),k}^{(i)}) \big) \Big\},$$

where $p(\theta_t | \theta_{t-1} = \hat{\theta}_{(t-1),k}^{(i)})$ is the probability distribution over the possible position of the ith particle with respect to the agent after it moves and is a Gaussian.

Simulation results validating the proposed search strategy for a team of three robots are shown Fig. 3. For a video of the simulation, we refer the reader to `https://youtu.be/5maQPeEcyf8`. The robots are shown localizing the source of a 2D multi-scale simulation of the Deepwater Horizon oil spill [3]. The entropy of the belief distribution for the source position at every time step is shown in Fig. 4a. We note that the temporary increase in entropy at 60^{th} time step corresponds to one of the robots loosing track of the plume.

3 Lagrangian Coherent Structures

While the proposed information theoretic search strategy described in the previous section is able to find and localize the source of a spill in a turbulent medium, the strategy can be significantly improved with GFD knowledge of the environmental dynamics. Coherent structures exhibited in GFD flows provide reduced-order description of the complex fluid environment and enable the estimation of the underlying geophysical fluid dynamics. In particular, Lagrangian coherent structures are important since they quantify transport and control the stretching and folding that underpins kinematic mixing.

Lagrangian coherent structures are material lines that organize fluid-flow transport and as mentioned may be viewed as the extensions of stable and unstable

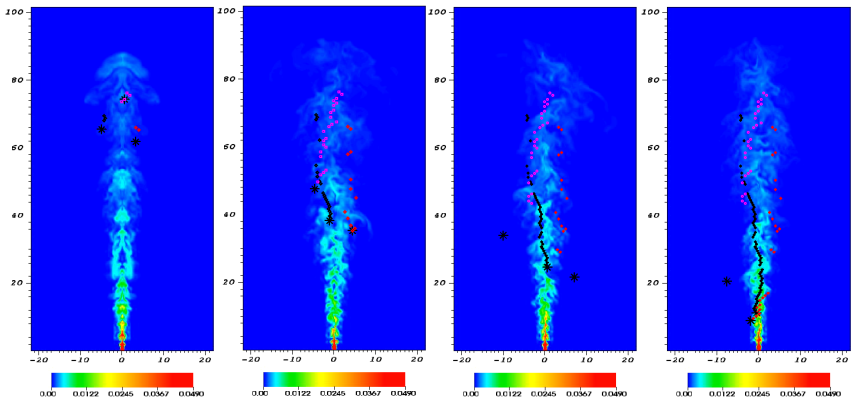

(a) (b) (c) (d)

Fig. 3: A team of three robots executing the proposed collaborative search strategy to localize the source of an oil spill. The background shows the gas volume concentration in the water due to the spill. The robots are denoted by black $*$. The red, magenta, and black dots denote the robot trajectories. For more details on the plume simulation, we refer the interested reader to [3].

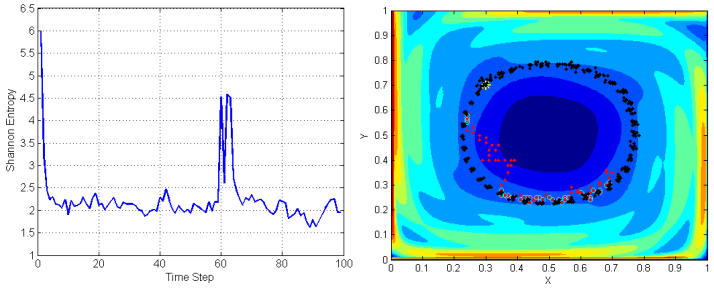

(a) (b)

Fig. 4: (a) The particles representing the hypotheses over the position of the source. (b) Single agent search using Algorithm 1 in a gyre flow. The background shows the FTLE field of the fluidic environment with LCS boundaries denoted in red. The source of the spill is denoted by the yellow $*$ and the black dots denote contaminants emanating from the yellow $*$. The red $*$ denotes the robot's current position, the red dots denote the robot's search trajectory. The cyan circles denote portions of the robot's trajectory where it detected the contaminants.

manifolds to general time-dependent systems [6]. In two-dimensional (2D) flows, LCS are one-dimensional separating boundaries analogous to ridges defined by local maximum instability, and can be quantified by local measures of Finite-Time Lyapunov Exponents (FTLE) [24, 5]. In this section, we briefly explain the computation of FTLE fields and the identification of LCS boundaries via local FTLE measures. We limit our discussions to 2D planar flows, however all concepts discussed in this section are readily extended to higher dimensions.

Consider a 2D flow field given by

$$\dot{\mathbf{x}}(t) = \mathbf{F}(\mathbf{x}, t) \tag{7}$$

where $\mathbf{x} = [x, y]^*$ gives the position in the plane and * denotes the transpose of the vector. Such a continuous time dynamical system has quantities, known as Lyapunov exponents, which are associated with the trajectory of the system in an infinite time limit. The Lyapunov exponents measure the growth rates of the linearized dynamics about the trajectory. To find the FTLE, one computes the Lyapunov exponents on a restricted finite time interval. For each initial condition, the exponents provide a measure of its sensitivity to small perturbations. Therefore, the FTLE is a measure of the local sensitivity to initial data. The position of a fluid particle advected by the flow field given by (7), is a function of time t, the starting point of the particle $\mathbf{x_0}$ and starting time t_0, *i.e.,* $\mathbf{x} = \mathbf{x}(t; \mathbf{x_0}, t_0)$. Using the notation by Shadden *et al.* [24], the solution to the dynamical system given in (7) can be viewed as a flow map which takes points from their initial position $\mathbf{x_0}$ at time t_0 to their position at time t. This map, denoted by $\phi_{t_0}^t$, satisfies $\phi_{t_0}^t(\mathbf{x_0}) = \mathbf{x}(t; \mathbf{x_0}, t_0)$, and has the properties: $\phi_{t_0}^{t_0}(\mathbf{x}) = \mathbf{x}$ and $\phi_{t_0}^{s+t}(\mathbf{x}) = \phi_s^{s+t}(\phi_{t_0}^s(\mathbf{x}))$.

The FTLE with a finite integration time interval T, associated with a point $\mathbf{x}$ at time t_0 is given by,

$$\sigma_{t_0}^T(\mathbf{x}) = \frac{1}{|T|} \ln \sqrt{\lambda_{max}(\Delta)} \tag{8}$$

where $\lambda_{max}(\Delta)$ is the maximum eigenvalue of the finite-time version of the Cauchy-Green deformation tensor Δ, given by,

$$\Delta = \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}^* \frac{d\phi_{t_0}^{t_0+T}(\mathbf{x})}{d\mathbf{x}}. \tag{9}$$

The value of Δ is computed numerically by discretizing the domain into a regular grid and computing the trajectories of each point and its immediate neighbors in the grid from time t_0 to $t_0 + T$. For each point in the grid, the trajectories are computed by numerically integrating (7) from t_0 to $t_0 + T$. The FTLE value gives a measure of the maximum expansion of two initially nearby particles when they are advected by the flow. Therefore, particles initiated on opposite sides of an LCS will have much higher FTLE values than their neighbors since LCS are essentially boundaries between two dynamically distinct regions in the flow.

By calculating the FTLE values for the entire flow field, it is possible to identify the LCS boundaries by tracing out regions with the highest FTLE values. As such, LCS are equivalent to ridges in the FTLE field with maximal FTLE values as defined by Shadden *et al.* [24]. The *forward-time* FTLE field calculated by advecting fluid particles forward in time ($T > 0$), reveals repelling LCS which are analogous to the stable manifolds of saddle points in a time independent flow field. Conversely, the *backward-time* FTLE field ($T < 0$) reveals attracting LCS which are analogous to unstable manifolds of a time independent flow field.

While one can easily extend the source seeking strategies described in Section 2 to single gyre-like flows as shown in Fig. 4b, knowing the LCS boundaries can significantly improve search strategies when operating in complex environments like those shown in Fig. 1a and 1b. Fig. 5 shows a simulation of the dispersion of particulates in a time-varying wind-driven double-gyre flow [31] with the FTLE field shown in the background. The LCS boundaries are denoted by red and the center vertical boundary oscillates about $x = 1$. From the simulations, we note that: 1) LCS boundaries behave as basin boundaries and thus fluid from opposing sides of the boundary do not mix; 2) in the presence of noise[2], particulates can cross the LCS boundaries and thus LCS denotes regions in the flow field where more escape events occur [4]; and 3) knowledge of LCS locations can improve the existing search strategy since this is akin to having a map of the workspace.

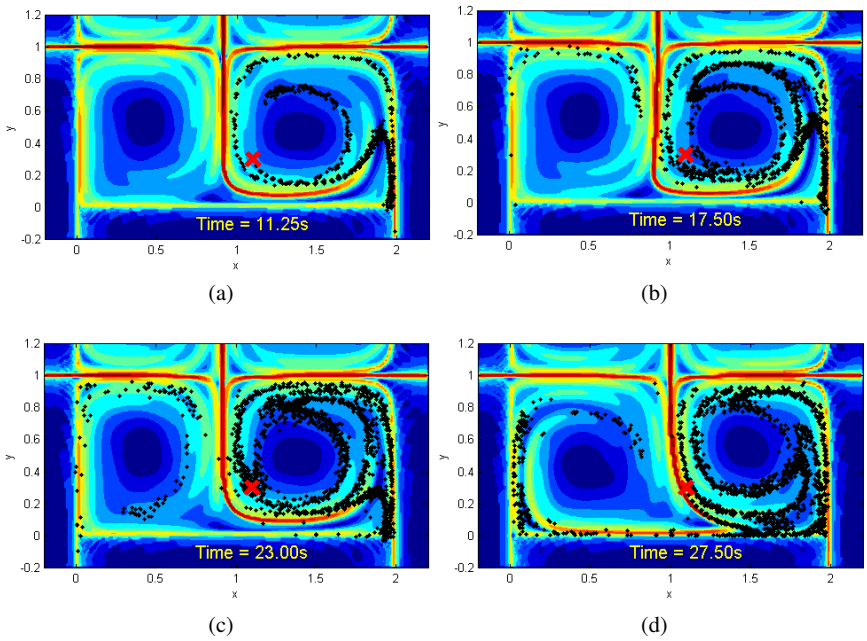

Fig. 5: Simulation of a spill in a time-varying wind-driven double-gyre flow. The background denotes the FTLE field for flow and the red x denotes the source position of the spill. The black particles denotes particulates emanating from the source.

However, one downside to finding attracting and repelling LCS by computing the FTLE fields is due to the fact that one needs global velocity field information for the region of interest. It is often the case when operating in the ocean that this information is sparse. Therefore, it would be useful to find the LCS in an alternate

[2] Noise can arise from uncertainty in model parameters and/or measurement noise.

way. To this end, we have developed a novel method that allows for collaborative robotic tracking of LCS. In particular, the robots perform the tracking based on local velocity measurements, thus ameliorating the need for global information. This collaborative LCS tracking is described in the following section.

4 Collaborative LCS Tracking

The tracking of coherent structures in fluids is challenging since they are global structures that are generally unstable and time-dependent. We briefly describe existing and on-going work in developing collaborative strategies for teams of robots to track the stable and unstable manifolds and LCS boundaries in 2D flows. It is important to note that tracking repelling LCS is achieved without relying on explicit computations of FTLE values. In fact, the identification of repelling LCS boundaries requires the computation of forward FTLEs, making real-time FTLE-based tracking of LCS boundaries extremely challenging. In contrast to this, tracking attracting LCS can be performed using backward FTLE fields calculated using the local velocity data acquired by the robots.

Repelling LCS and Stable Manifolds The tracking of repelling LCS boundaries or stable manifolds in flows relies on robots maintaining a boundary straddling formation while collecting local measurements of the flow field. Since LCS correspond to regions in the flow field with extremal velocities, LCS tracking is achieved by fusing data obtained on opposite sides of the boundary to identify the location of the extremal velocity. The strategy was first proposed for a team of three robots where the center robot is responsible for locating the boundary using the flow measurements provided by the team. The remaining robots would then maintain a saddle straddling formation, *i.e.*, remain on opposite sides of the boundary, at all times. The strategy has been extensively validated using analytical models, experimental data, and actual ocean data [9, 19, 18] and extended for larger team sizes [20].

It is important to note that these strategies achieve tracking of *global* fluidic features, *i.e.,* the LCS, using only *local* measurements of the flow field and an initial estimate of the location of the repelling LCS.

Attracting LCS and Unstable Manifolds Different from repelling LCS, attracting LCS are quantified by maximal backward FTLE measures. As such, collaborative tracking of attracting LCS or unstable manifolds can be achieved through onboard calculation of local FLTE fields using previously acquired flow velocity data. In this case, the tracking strategy utilizes the FTLE field along with instantaneous local flow field measurements to resolve the attracting LCS boundary. Similar to tracking repelling LCS, agent-level control policies leveraging underlying flow dynamics were developed to maintain the formation of the team as they track the attracting boundary. The formation control strategy ensures vehicles do not collide with one another while maintaining the necessary boundary straddling formation.

The attracting LCS tracking strategy was also validated using a combination of analytical models and experimental flow tanks using micro autonomous surface ve-

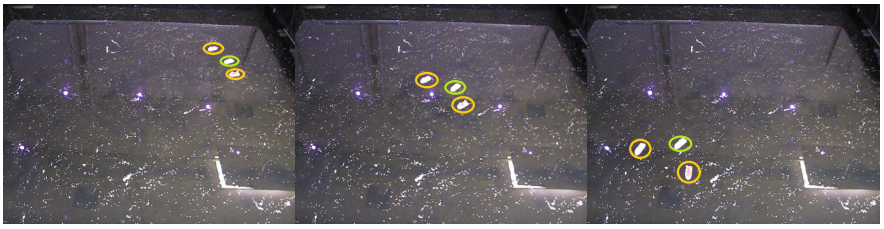

(a) t=9.3s (b) t=20.4s (c) t=31.6s

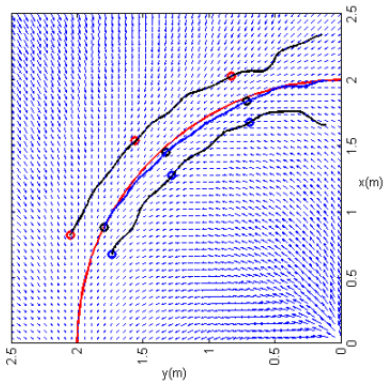

(d)

Fig. 6: (a)-(c) Actual ASV tracking experiment. (d) ASVs trajectories for the experiment in (a)-(c). The red line is the actual boundary and the blue line is the trajectory of the center vehicle.

hicles (ASVs) [14]. Fig. 6 shows an experiment using three real and four virtual ASVs to track a simulated static flow field. The three ASVs were initially arranged in a saddle straddling formation with the center ASV tasked with tracking the boundary. The virtual agents were placed at the four corners of the grid.

Collaborative tracking of coherent structures allows one to gain knowledge of the geophysical flow of interest. This knowledge can be used to improve or optimize a variety of sensing and control objectives. For example, in previous work we have demonstrated how GFD-based knowledge allows for an increase in loitering time of vehicles operating in the ocean [4, 17, 10, 7]. Another example involves the localization of a contaminant source as described previously.

5 Conclusions and Future Outlook

In this work, we showed how GFD-derived knowledge can significantly improve the autonomy of the vehicles that operate within them. Coherent structures have the potential to provide computationally efficient forecasts of current and wind patterns

since they enable a much lower order description of the environmental dynamics. The ability to better quantify transport behaviors in natural fluid environments will enable the synthesis of energy-efficient motion planning and control strategies, the detection and tracking of contaminant and hazardous waste dispersions, and more effective allocation of mobile sensing resources for search and rescue operations. This makes intuitive sense since this is akin to planning and motion control of autonomous vehicles using a suitable map of the environment.

Different from traditional maps, most features of interest in flows are unstable and non-stationary. While this renders the problem of leveraging dynamic transport controlling structures for improved underwater vehicle autonomy highly challenging, it also opens up new opportunities in planning, control, and perception in these highly dynamic and uncertain environments. Some specific directions for future investigations include how do we construct, maintain, and update such a fluidic map? Since LCS boundaries exist at various length scales and are often seasonal, how does a vehicle leverage the structures for energy efficient monitoring and navigation? Is it possible to use LCS information to identify unsafe regions in the fluidic environment where vehicles may be trapped by the environmental dynamics? Can we manipulate the fluid for fluid-mediated underwater manipulation of objects?

Beyond these challenging problems is the extension of all the previously discussed tracking, planning, control, search, etc. problems to the fully 3D environment. One major challenge associated with autonomous vehicles operating at depth are the issues involved with communication. It is known that communication time delay can destabilize formations of vehicle groups [29]. But one also must consider communication drop-outs and even the complete loss of communication especially when operating in the ocean. Overcoming the difficulties of operating in a 3D environment will allow for greatly improved environmental monitoring and forecasting.

In this work, we presented information theoretic search strategies for localizing contaminant sources in turbulent flows. We also presented on-going work in distributed sensing of geophysical fluid dynamics. Much of this work has been motivated by the insight that the environment may be effectively exploited for vehicle control and autonomy especially when the vehicle and environmental dynamics are tightly coupled. By looking at the changing environment through a geophysical perspective, there are significant GFD features that can be leveraged for predicting and estimating the environmental dynamics. The challenge lies in overcoming the theoretical and technological challenges needed to robustly and autonomously collect, process, and interpret data about the geophysical flows.

Acknowledgements

MAH, HH, and DK are supported by ONR Award Nos. N000141211019 and N000141310731 and National Science Foundation (NSF) grant IIS-1253917. EF and PAY are supported by National Science Foundation (NSF) grant DMS-1418956. IBS is supported by ONR contract No. N0001412WX2003 and NRL 6.1 program

contract No. N0001412WX30002. We would like to thank Alex Fabregat Tomas (CUNY) and Andrew Poje (CUNY) in providing the plume data.

References

[1] Charrow B, Michael N, Kumar V (2013) Cooperative Multi-robot Estimation and Control for Radio Source Localization. Springer International Publishing
[2] Crisan D, Doucet A (2002) A survey of convergence results on particle filtering methods for practitioners. IEEE Trans on Signal Processing 50(3):736–746
[3] Fabregat A, Dewar WK, Ozgokmen TM, Poje AC, Wienders N (2015) Numerical simulations of turbulent thermal, bubble and hybrid plumes. Ocean Modelling 90:16–28
[4] Forgoston E, Billings L, Yecko P, Schwartz IB (2011) Set-based corral control in stochastic dynamical systems: Making almost invariant sets more invariant. Chaos 21(013116)
[5] Haller G (2011) A variational theory of hyperbolic Lagrangian coherent structures. Physica D 240:574–598
[6] Haller G, Yuan G (2000) Lagrangian coherent structures and mixing in two-dimensional turbulence. Phys D 147:352–370
[7] Heckman CR, Hsieh MA, Schwartz IB (2014) Controlling basin breakout for robots operating in uncertain flow environments. In: International Symposium on Experimental Robotics (ISER 2014), Marrakech/Essaouira, Morocco
[8] Hoffmann G, Tomlin C (2010) Mobile sensor network control using mutual information methods and particle filters. Automatic Control, IEEE Transactions on 55(1):32–47
[9] Hsieh MA, Forgoston E, Mather TW, Schwartz I (2012) Robotic manifold tracking of coherent structures in flows. In: Proc. IEEE Int. Conf. on Robotics and Automation (ICRA2012), Minneapolis, MN USA
[10] Hsieh MA, Mallory K, Schwartz IB (2014) Distributed allocation of mobile sensing agents in geophysical flows. In: Proc. of the 2014 American Controls Conference, Portland, OR
[11] Huber M, Bailey T, Durrant-Whyte H, Hanebeck U (2008) On entropy approximation for gaussian mixture random vectors. In: Multisensor Fusion and Integration for Intelligent Systems, 2008. MFI 2008. IEEE International Conference on, pp 181–188, DOI 10.1109/MFI.2008.4648062
[12] I TG (1953) Dispersion of soluble matter in solvent flowing slowly through a tube. Proc R Soc Lond A A219:186–203
[13] Inanc T, Shadden S, Marsden J (2005) Optimal trajectory generation in ocean flows. In: Proc. of the American Control Conference, pp 674 – 679
[14] Kularatne D, Hsieh A (2015) Tracking attracting lagrangian coherent structures in flows. In: Proceedings of Robotics: Science and Systems, Rome, Italy
[15] Levy M, Ferrari R, Franks PJS, Martin AP, Riviere P (2012) Bringing physics to life at the submesoscale. Geophysical Research Letters 39(14)

[16] Lolla T, Ueckermann MP, Haley P, Lermusiaux PFJ (2012) Path planning in time dependent flow fields using level set methods. In: Proc. IEEE International Conference on Robotics and Automation, Minneapolis, MN USA

[17] Mallory K, Hsieh MA, Forgoston E, Schwartz IB (2013) Distributed allocation of mobile sensing swarms in gyre flows. Nonlin Processes Geophys 20(5):657–668

[18] Michini M, Hsieh MA, Forgoston E, Schwartz IB (2014) Experimental validation of robotic manifold tracking in gyre-like flows. In: Proc. IEEE/RSJ Int. Conf. on Intelligent Robots and Systems (IROS) 2014, Chicago, IL USA

[19] Michini M, Hsieh MA, Forgoston E, Schwartz IB (2014) Robotic tracking of coherent structures in flows. IEEE Trans on Robotics 30(3):593–603

[20] Michini M, Rastgoftar H, Hsieh MA, Jayasuriya S (2014) Distributed formation control for collaborative tracking of manifolds in flows. In: Proc. of the 2014 American Control Conference (ACC 2014), Portland, OR

[21] Russell R (2004) Locating underground chemical sources by tracking chemical gradients in 3 dimensions. In: Intelligent Robots and Systems, 2004. (IROS 2004). Proceedings. 2004 IEEE/RSJ International Conference on, vol 1, pp 325–330 vol.1, DOI 10.1109/IROS.2004.1389372

[22] SCRIPPS (2014) Naitonal HF RADAR network - surface currents. URL `http://cordc.ucsd.edu/projects/mapping/maps/`

[23] Senatore C, Ross S (2008) Fuel-efficient navigation in complex flows. In: American Control Conference, 2008, pp 1244 –1248

[24] Shadden SC, Lekien F, Marsden JE (2005) Definition and properties of Lagrangian coherent structures from finite-time Lyapunov exponents in two-dimensional aperiodic flows. Phys D: Nonlin Phenomena 212(3-4):271 – 304

[25] Shchepetkin A, McWilliams J (1998) Quasi-monotone advection schemes based on explicit locally adaptive dissipation. Monthly Weather Review 126:1541–1580

[26] Shchepetkin AF, McWilliams JC (2005) The regional oceanic modeling system (ROMS): a split-explicit, free-surface, topography-following-coordinate oceanic model. Ocean Modeling 9:347–404

[27] Smith R, Kelly J, Sukhatme G (2012) Towards improving mission execution for autonomous gliders with an ocean model and kalman filter. In: Proc. of the 2012 IEEE Int. Conf. on Robotics and Automation, Minneapolis, MN

[28] Smith RN, Chao Y, Li PP, Caron DA, Jones BH, Sukhatme GS (2010) Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model. International Journal of Robotics Research 29(12):1475–1497

[29] Mier-y Teran-Romero L, Forgoston E, Schwartz I (2012) Coherent pattern prediction in swarms of delay-coupled agents. IEEE Trans on Robotics 28(5):1034–1044

[30] Vergassola M, Villermaux E, Shraiman BI (2007) Infotaxis as a strategy for searching without gradients. Nature 445:406–409

[31] Veronis G (1966) Wind-driven ocean circulation, part i and part ii. Deep-Sea Res 13(31)