Active Perception for Plume Source Localisation with Underwater Gliders

Ki Myung Brian Lee[1], James Ju Heon Lee[1], Chanyeol Yoo[1], Ben Hollings[2] and Robert Fitch[1]
[1] Centre for Autonomous Systems, University of Technology Sydney, NSW Australia
{brian.lee, juheon.lee}@student.uts.edu.au, {chanyeol.yoo, robert.fitch}@uts.edu.au
[2] Blue Ocean Monitoring, WA Australia
ben.hollings@blueoceanmonitoring.com.au

Abstract

We consider the problem of localising an unknown underwater plume source in an energy-optimal manner. We first develop a specialised Gaussian process (GP) regression technique for estimating the source location given concentration measurements and an ambient flow field. Then, we use the GP upper confidence bound (GP-UCB) for active perception to choose sampling locations that both improve the estimate of the source and lead the glider to the correct source location. A trim-based FMT*planner is then used to find the sequence of controls that minimise the energy consumption. We provide a theoretical guarantee on the performance of the algorithm, and demonstrate the algorithm using both artificial and experimental datasets.

1 Introduction

Plume source localisation is a prominent class of algorithmic problems with immediate and critical practical applications in environmental monitoring [Rudnick *et al.*, 2004], chemical warfare protection [Gunatilaka *et al.*, 2008], and oil and gas source localisation [Russell-Cargill *et al.*, 2018]. We are interested in the underwater context, where the goal is to find the source of a substance that is subject to an ambient flow field.

There are various circumstances in which it would be beneficial to track substances diffusing in water. One interesting recent development is a sensor that can measure methane concentration and that is suitable for use with underwater robots [Russell-Cargill *et al.*, 2018]. It is imperative to develop algorithms that make use of such sensors for effective application over large spatial scales of tens to hundreds of kilometres. We treat this problem in the active perception context, and consider how to find the source of a plume while improving its estimate simultaneously.

A major challenge is how to take into account realistic ocean current models. Most existing work assumes uniform current, which is restrictive. Relaxing this assumption is difficult, however, because the relationship between plume source and concentration with varying ocean current is described by a nonlinear partial differential equation (PDE), and cannot be written in a general closed form.

We address this challenge by elegantly incorporating the PDE relationship into Gaussian process (GP) regression. We first use a GP to represent the concentration, and then use the PDE model to relate the concentration to the source of the plume. Then, we apply an existing algorithm (GP-UCB) to perform active search for the source.

In this paper, we present a novel algorithm for general plume source localisation supported by rigorous theoretical analysis. We evaluate the algorithm using artificial data, and also using real data previously collected by an underwater glider equipped with a methane concentration detector over a period of two weeks [Russell-Cargill *et al.*, 2018]. Results in simulation show that our algorithm can accurately find the ground truth source of an artificial plume. We cannot make the same statement for the case with real data, as the true source location is unknown, but results are consistent with intuition.

Our algorithm is general and can be used for many different types of plumes. This is an exciting result that enables applications of underwater gliders and other types of AUVs that cannot currently be addressed in any other way. Application to methane plumes in particular is significant because methane is used as an indicator to directs gas and oil exploration, which is an activity with substantial economic value.

2 Related Work

2.1 Plume Source Localisation

Autonomous monitoring of environmental plumes is of substantial interest [Kowadlo and Russell, 2008]. There are two main relevant bodies of work. The first is concerned with tracking the boundary of a plume, and the other is finding the source of a plume. Tracking the

boundary of a plume is useful for scenarios such as oil leaks in the ocean, and has been achieved by using heuristic strategies [Smith *et al.*, 2010], level set-based tracking [Marthaler and Bertozzi, 2004], and by exploiting the structure of a flow field [Kularatne *et al.*, 2015]. Multi-robot algorithms have also been proposed to address the large spatial extent of plumes [Li *et al.*, 2014].

In this work, we are interested in the problem of locating the source of plume. A traditional approach is to adopt an empirical parametric model referred to as the Gaussian plume [Gunatilaka *et al.*, 2008], and to fit the observations to the model. However, using a parametric model may prove demanding in practical scenarios, because the data can be sparse or subject to noise. Adaptive strategies such as waiting time [Chang *et al.*, 2013] and centroid tracking [Smith *et al.*, 2010] have been proposed that improve their accuracy given additional data. A more principled approach is the 'infotaxi' strategy [Vergassola *et al.*, 2007], where actions are chosen to maximise the information gain over the source location [Hajieghrary *et al.*, 2017]. Although the infotaxi strategy has a theoretically justified background, the probability distribution used for calculating the information gain is not given in closed form for any flow fields more complex than the constant case, which is a substantial limitations.

The algorithm we propose in this paper is more flexible in terms of the types of flow fields and plume sources. We incorporate a physical model of plume dispersion into the Gaussian process (GP) regression framework, which allows us to estimate the source location in a *nonparametric* manner. Further, the tracking strategy proposed in this paper has a known theoretical performance guarantee which, as we show, still holds in our case.

2.2 Motion Planning for Underwater Gliders

Autonomous underwater gliders are a class of autonomous underwater vehicles (AUVs). They are distinguished from other AUVs by their propulsion mechanism, where the forward motion is generated by changing its buoyancy to dive up and down. The lack of an active propulsion makes underwater gliders more energy efficient compared to other AUVs at the expense of speed, making them sensitive to the surrounding ocean flow field. These qualities makes gliders ideal for long duration missions such as surveying and patrolling, as long as it can properly exploit the ocean current.

In this work, we are interested in finding an energy-optimal 3D glider path in an ocean flow field. Popular planning algorithms such as RRT [Rao and Williams, 2009], A* [Fernández-Perdomo *et al.*, 2010; Isern-González *et al.*, 2011; Fernández-Perdomo *et al.*, 2011; Xinke *et al.*, 2015], and genetic algorithms [Zamuda and Hernández Sosa, 2014; Shih *et al.*, 2017] have been used as planning frameworks to properly navigate the underwater glider through the 2D ocean flow field. Underwater glider manoeuvres in 3D space also exists, where Dubins path made with a series of 'sawtooth' and spiral motions is given as a solution [Cao *et al.*, 2017; Liu *et al.*, 2017]. However these work do not consider the ocean flow field. Also in our recent work on underwater glider planning [Lee *et al.*, 2017], we have demonstrated that traditional operations of underwater gliders may not give the most energy-optimal path. Most work assumes directly controllable turning rates and simplified kinematics model that consequently neglects important glider states, such as glider angle and the ballast.

In this work, we will be using the trim-based FMT*planning algorithm demonstrated in [Lee *et al.*, 2017] to find the energy-optimal glider path trajectory in 3D ocean environment. The planner makes use of 'trim-states' that holds the critical glider states needed for accurate glider model.

3 Background

3.1 Underwater Glider

Underwater gliders makes forward velocity by using the lift and drag force generated from the sequences of dives. The forces are dependent on the glider's angle of attack that is a function of its glide angle.

A continuous dynamical model of an underwater glider with high dimensions can be expressed as follows:

$$\dot{\mathbf{x}}(t) = f(\mathbf{x}(t), \mathbf{u}(t)) + \mathbf{w}(\mathbf{x}(t)), \tag{1}$$

where $\mathbf{x}(t)$ is the glider state at time t, $\mathbf{u}(t)$ is the control vector, and $\mathbf{w}$ is the ocean current vector at current glider position given *a priori*. The control vector consists of; the ballast pump control used to regulate the glider's buoyancy for dives and surfacing, and the moving mass position used to control glider's orientation and velocity, as velocity is a function of the glide angle.

A *trim-state* is a state of dynamic equilibrium that a vehicle will hold to in the absence of disturbances or variations to control [Leonard and Graver, 2001]. As we expect the glider to travel long distances between each state, we can reformulate the dynamic model as a trim-state, where the time taken to change the control state is assumed negligible. Since no energy is spent while travelling between states, the trim-state can be reduced down to a kinematic model that considers the control vector:

$$\begin{bmatrix} V_k & \gamma_k & \delta_k & m_{bk} \end{bmatrix}^T, \tag{2}$$

where V_k is the speed, γ is the glide angle, δ is the heading angle, and m_b is the ballast mass (i.e., amount of water in the ballast tank). We assume that the ballast

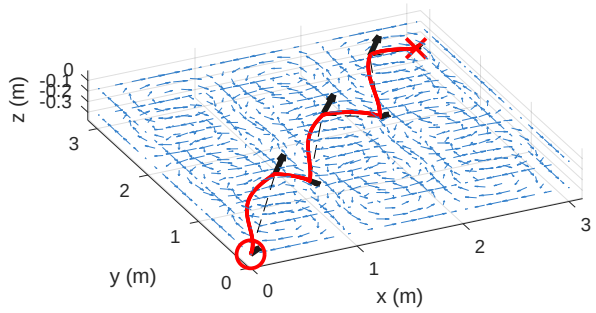

Figure 1: Trim-based glider manoeuvre over 3D environment under the influence of ocean currents

tank is either empty or full (i.e., $m_b \in [0, m_{b\,\max}]$). The new glider kinematic motion can be described as:

$$\mathbf{x}(t+1) = \mathbf{x}(t) + \left(\begin{bmatrix} V_k(\gamma)\cos\gamma\cos\delta \\ V_k(\gamma)\cos\gamma\sin\delta \\ V_k(\gamma)\sin\gamma \end{bmatrix} + \mathbf{w}(\mathbf{x}(t)) \right) \Delta t. \tag{3}$$

and use it to solve for the trim-state that moves the glider from one state to another. Given a sequence of trim states, an example of a continuous glider trajectory under the influence of ocean currents is shown in Fig. 1.

The underwater glider energy cost is the sum of the ballast pump cost, moving mass re-position cost, turning cost, and the hotel cost. Note that the pump cost is only applied when expelling water out of the ballast tank (i.e. when increasing buoyancy), since the static pressure from the ocean environment can be used to fill the tank without the pump.

3.2 Plume Dispersion Model

Advection-diffusion equation [Lupini and Tirabassi, 1979] describes the relation between source and plume under the influence of flow field. Advection refers to the transport of plume particles by the flow field, and diffusion describes the natural dispersion in still water.

The introduction of plume species into the environment by a scalar function is $s : \mathbb{R}^+ \times \mathbb{R}^3 \to \mathbb{R}$, referred to as the *source strength*. The source strength describes the influx of species into the environment (i.e. a source) at position $\mathbf{x}$ and time t. Given the source strength, $s(\mathbf{x}, t)$, the advection-diffusion equation describes the evolution of concentration of a species, $c : \mathbb{R}^+ \times \mathbb{R}^3$ as:

$$\frac{\partial c(\mathbf{x},t)}{\partial t} + \mathbf{w}(\mathbf{x},t) \cdot \nabla c(\mathbf{x},t) = D\Delta^2 c(\mathbf{x},t) + s(\mathbf{x},t), \tag{4}$$

where ∇ is the gradient operator, Δ^2 is the Laplacian operator, $\mathbf{w}$ is the flow field and D is the diffusion coefficient. For instance, the diffusion coefficient for methane-water system ranges between $D = 1 - 5m^2/s$, depending on the temperature [Mrazovac *et al.*, 2012].

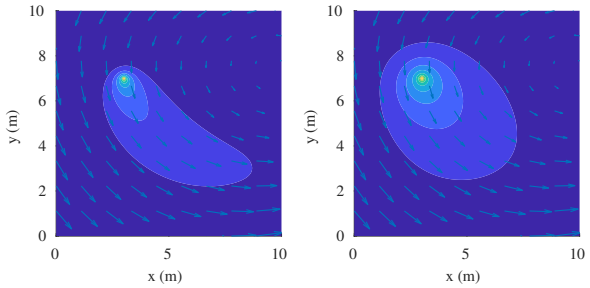

(a) Advection-dominant (b) Diffusion-dominant

Figure 2: Examples of 2D advection-diffusion process

Given $s(\mathbf{x})$ is time invariant, the plume process reaches its steady-state. Based on (4), the steady-state equation is obtained by setting the time derivative to zero:

$$\mathbf{w}(\mathbf{x}) \cdot \nabla c(\mathbf{x}) = D\Delta^2 c(\mathbf{x}) + s(\mathbf{x}). \tag{5}$$

where we wrote $c(\mathbf{x}) = \lim_{t\to\infty} c(\mathbf{x}, t)$ with a slight abuse of notation.

If $\mathbf{w}(\mathbf{x})$ does not vary across space, it is well-known that (5) admits the so-called Gaussian plume solution [Lupini and Tirabassi, 1979]. However, it is practically impossible to find an analytical solution in the general case when $\mathbf{w}(\mathbf{x})$ varies spatially, and (5) needs to be solved numerically using, e.g. finite difference method. An example of such a solution in 2D is shown in Fig. 2. In Fig. 2, it can be seen that the chemical species is not only advected along the flow, but also diffused in directions against the flow.

4 Problem Statement

Suppose we have an underwater glider described in Sec. 3.1 with trim state-based manoeuvre under the influence of fully known, spatially varying ocean currents $\mathbf{w}$. There is a plume underwater of which the concentration c is governed by the dispersion model described in Sec. 3.2. We will model the source of the plume by *source strength*, $s(\mathbf{x})$, which describes the amount of plume species being introduced to the environment at position $\mathbf{x}$. Given the source strength, we model the source location by:

$$\mathbf{x}_s = \arg\max_{\mathbf{x}\in\mathcal{E}} s(\mathbf{x}), \tag{6}$$

where $\mathcal{E}$ is the operating environment of the glider.

When the glider is near the seabed, it measures the concentration of the plume at a point using an onboard sensor. The measurements are modelled by:

$$y = c(\mathbf{x}) + \epsilon, \tag{7}$$

where $\epsilon \sim \mathcal{N}(0, \sigma_y^2)$ is white Gaussian noise with known variance σ_y^2. We will denote a set of observations as $\mathcal{O} = \{\mathbf{x}_i, y_i\}_{i=1}^N$.

Given the concentration measurements, we wish to find a sequence of optimal control vectors $\mathbf{u}(t)$ that leads to the source location $\mathbf{x}_s$ in an energy-optimal manner. Formally, we wish to find a sequence of controls, $\mathbf{u}(t)$ that solves the following problem.

$$\begin{aligned} \underset{u(t)}{\text{minimise}} \quad & \sum_{t=1}^{T} E(\mathbf{x}(t), \mathbf{u}(t)) \\ \text{subject to} \quad & \dot{\mathbf{x}} = \mathbf{f}(\mathbf{x}, \mathbf{u}) \\ & \mathbf{x}(T) = \mathbf{x}_s. \end{aligned} \tag{8}$$

where $E(\mathbf{x}(t), \mathbf{u}(t))$ is the energy cost function.

As the source strength over the environment is initially unknown, the overall problem necessitates two main sub-problems: *estimation* and *planning*. The estimation problem is to predict the source strength over the environment given sparse measurements of the plume concentration. To do so, we first estimate the plume concentration over the environment, and predict the source strength by exploiting a physical model. An overview of the framework is illustrated in Fig. 3.

Based on the uncertain estimate of the source strength, the planning problem is to find a sequence of controls that is energy-optimal and leads the glider to the source in a *probabilistic sense*. Note that because the true source strength is unknown, and the estimation of the true source strength also depends on the actions taken, the planner must also consider the uncertainty of the estimated source strength, and encourage increase in information gain.

It is difficult to directly solve for a sequence of controls that optimises the energy expenditure while actively pursuing the source location, because it implicitly requires maximising the unknown source strength function. Thus, we solve the planning problem in a hierarchical manner, consisting of a dive-location planner and a trajectory planner. The dive-location planner chooses dive locations that leads the glider to the goal, while actively improving the estimate of the source. This is achieved by using GP upper confidence bound (GP-UCB) strategy. Based on the dive-locations generated, the trajectory planner plans an energy-optimal path through these dive-locations using the FMT*framework.

5 GP-Based Source Localisation

The estimation part of the framework estimates the source location by estimating the source strength given concentration measurements. It consists of two components: *concentration estimation* and *source localisation*.

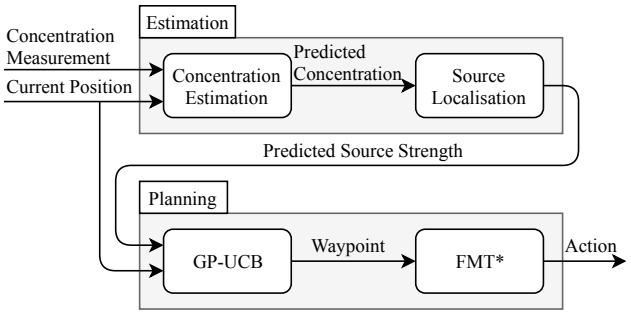

Figure 3: Overview of hierarchical framework

5.1 Estimation of Plume Concentration

The plume concentration at point $\mathbf{x}$ modelled as a GP is denoted as $c(\mathbf{x})$. Formally,

$$c(\mathbf{x}) \sim GP(0, k_{cc}(\mathbf{x}, \mathbf{x}')), \tag{9}$$

where $k_{cc}(\mathbf{x}, \mathbf{x}')$ is a kernel function. Since the fundamental solutions to (5) is given by a Gaussian plume, we use the squared-exponential (SE) kernel:

$$k_{SE}(\mathbf{x}, \mathbf{x}') = \sigma_c^2 \exp\left(-\frac{\gamma}{2}\|\mathbf{x} - \mathbf{x}'\|^2\right), \tag{10}$$

where σ_c^2 is a self-variance and γ is the length-scale parameter.

Given a set of observations $\mathcal{O} = \{\mathbf{x}_i, y_i\}_{i=1}^N$ in the form of (7), we learn the GP-based concentration function $c(\mathbf{x}$ by representing observations and concentration as a joint random variable. Given the joint random variable of the form $\mathbf{y}_\mathcal{O} = \begin{bmatrix} y_1 & \cdots y_N \end{bmatrix}^T$, we have

$$\begin{bmatrix} \mathbf{y}_\mathcal{O} \\ c(\mathbf{x}^*) \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K_{cc} + \sigma_y^2 I & k^* \\ k^{*T} & k^{**} \end{bmatrix}\right) \tag{11}$$

where, $k_{cc}^{**} \in \mathbb{R}$, $\mathbf{k}_{cc}^* \in \mathbb{R}^N$, and $K_{cc} \in \mathbb{R}^{N \times N}$ are:

$$\begin{aligned} k_{cc}^{**} &= k_{cc}(\mathbf{x}^*, \mathbf{x}^*) \\ \mathbf{k}_{cc}^* &= \begin{bmatrix} k_{cc}(\mathbf{x}_1, \mathbf{x}^*) & \cdots & k_{cc}(\mathbf{x}_N, \mathbf{x}^*) \end{bmatrix}^T \\ K_{cc} &= \begin{bmatrix} k_{cc}(\mathbf{x}_1, \mathbf{x}_1) & \cdots & k_{cc}(\mathbf{x}_1, \mathbf{x}_N) \\ \vdots & \ddots & \vdots \\ k_{cc}(\mathbf{x}_N, \mathbf{x}_1) & \cdots & k_{cc}(\mathbf{x}_N, \mathbf{x}_N) \end{bmatrix}. \end{aligned} \tag{12}$$

Based on the joint random variable, the conditional distribution of the concentration $c(\mathbf{x})$ at the query point $\mathbf{x}$ given a set of observations $\mathcal{O}$ is described as:

$$\mathcal{P}(c(\mathbf{x}^*) \mid \mathcal{O}) = \mathcal{N}(\mu_c(\mathbf{x}^*), \sigma_c^2(\mathbf{x}^*)), \tag{13}$$

where the conditional mean μ and variance σ^2 are

$$\begin{aligned} \mu(\mathbf{x}^*) &= \mathbf{k}_{cc}^{*T}(K_{cc} + \sigma_y^2 I)^{-1}\mathbf{y}_\mathcal{O} \\ \sigma^2(\mathbf{x}^*) &= k_{cc}^{**} - \mathbf{k}_{cc}^{*T}(K_{cc} + \sigma_y^2 I)^{-1}\mathbf{k}_{cc}^*. \end{aligned} \tag{14}$$

5.2 Localisation of Plume Source

Given the concentration $c(\mathbf{x})$ estimated using GP regression in (13), we exploit the advection-diffusion PDE (5) to estimate the source strength. In order to do so, we need to combine the PDE into the GP regression framework. This can be achieved by considering the PDE as an operator on the vector space of the function [Särkkä and Hartikainen, 2012]. The advection-diffusion PDE (5) can be written in an operator-theoretic form by introducing:

$$\mathcal{D}_{\mathbf{x}} \equiv (\mathbf{w}(\mathbf{x}) \cdot \nabla - D\Delta^2). \tag{15}$$

By rearranging (5), we find that $s(\mathbf{x}) = \mathcal{D}_{\mathbf{x}} c(\mathbf{x})$. In other words, we can recover the source strength from the concentration by applying $\mathcal{D}_{\mathbf{x}}$ on $c(\mathbf{x})$. Using this property of $\mathcal{D}_{\mathbf{x}}$, it can be shown that the kernel function for $s(\mathbf{x})$ and the cross-covariance functions are given by applying $\mathcal{D}_{\mathbf{x}}$ as:

$$k_{sc}(\mathbf{x}, \mathbf{x}') = Cov(s(\mathbf{x}), c(\mathbf{x}')) = \mathcal{D}_{\mathbf{x}} k_{cc}(\mathbf{x}, \mathbf{x}'), \tag{16}$$
$$k_{cs}(\mathbf{x}, \mathbf{x}') = Cov(c(\mathbf{x}), s(\mathbf{x}')) = k_{cc}(\mathbf{x}, \mathbf{x}') \mathcal{D}_{\mathbf{x}'}, \tag{17}$$
$$k_{ss}(\mathbf{x}, \mathbf{x}') = Cov(s(\mathbf{x}), s(\mathbf{x}')) = \mathcal{D}_{\mathbf{x}} k_{cc}(\mathbf{x}, \mathbf{x}') \mathcal{D}_{\mathbf{x}'}. \tag{18}$$

It is important to note that (16), (17) and (18) can be calculated analytically for a given choice of kernel $k_{cc}(\mathbf{x}, \mathbf{x}')$ for concentration. For instance, assuming SE kernel (10), the cross-covariance function is given by:

$$\begin{aligned} &k_{sc}(\mathbf{x}, \mathbf{x}') \\ =&(D(2\gamma - \gamma^2 \|\mathbf{x}\|^2) + \gamma \mathbf{w}(\mathbf{x})^T (\mathbf{x} - \mathbf{x}')) k_{SE}(\mathbf{x}, \mathbf{x}'). \end{aligned} \tag{19}$$

Using the cross-covariance kernel functions (16), (17) and (18), we can rewrite the concentration $c(\mathbf{x})$ and source strength $s(\mathbf{x})$ as a joint GP [Särkkä and Hartikainen, 2012]:

$$\begin{bmatrix} c(\mathbf{x}) \\ s(\mathbf{x}) \end{bmatrix} \sim GP\left(\mathbf{0}, \begin{bmatrix} k_{cc}(\mathbf{x}, \mathbf{x}') & k_{cs}(\mathbf{x}, \mathbf{x}') \\ k_{sc}(\mathbf{x}, \mathbf{x}') & k_{ss}(\mathbf{x}, \mathbf{x}') \end{bmatrix}\right). \tag{20}$$

Combining (11) with (20), we get a joint random variable for observation $\mathbf{y}_{\mathcal{O}}$ and the source strength $s(\mathbf{x}^*)$ at a query point $\mathbf{x}^*$ denoted as

$$\begin{bmatrix} \mathbf{y}_{\mathcal{O}} \\ s(\mathbf{x}^*) \end{bmatrix} \sim \mathcal{N}\left(\mathbf{0}, \begin{bmatrix} K_{cc} + \sigma_y^2 I & \mathbf{k}_{cs}^* \\ \mathbf{k}_{sc}^{*T} & k_{ss}^{**} \end{bmatrix}\right). \tag{21}$$

where $k_{ss}^{**} \in \mathbb{R}$, $\mathbf{k}_{sc}^* \in \mathbb{R}^N$, and $\mathbf{k}_{cs}^* \in \mathbb{R}^N$ are given by:

$$\begin{aligned} k_{ss}^{**} &= k_{ss}(\mathbf{x}^*, \mathbf{x}^*) \\ \mathbf{k}_{sc}^* &= \begin{bmatrix} k_{sc}(\mathbf{x}^*, \mathbf{x}_1) & \dots & k_{sc}(\mathbf{x}^*, \mathbf{x}_N) \end{bmatrix}^T \\ \mathbf{k}_{cs}^* &= \begin{bmatrix} k_{cs}(\mathbf{x}_1, \mathbf{x}^*) & \dots & k_{cs}(\mathbf{x}_N, \mathbf{x}^*) \end{bmatrix}^T. \end{aligned}$$

The source strength at the query point can then be predicted as the conditional distribution [Rasmussen and Williams, 2006]:

$$\mathcal{P}(s(\mathbf{x}^*) \mid \mathcal{O}) = \mathcal{N}(\mu_c(\mathbf{x}^*), \sigma_c^2(\mathbf{x}^*)), \tag{22}$$

where the mean and variance are

$$\begin{aligned} \mu(\mathbf{x}^*) &= \mathbf{k}_{sc}^{*T} (K_{cc} + \sigma_y^2 I)^{-1} \mathbf{y}_{\mathcal{O}} \\ \sigma^2(\mathbf{x}^*) &= k_{ss}^{**} - \mathbf{k}_{sc}^{*T} (K_{cc} + \sigma_y^2 I)^{-1} \mathbf{k}_{cs}^*. \end{aligned} \tag{23}$$

6 Hierarchical Planning for Source Seeking

Given the GP model of source strength in Sec. 5, we find control vectors for finding the plume source. We first find the most probable next dive-to waypoint given the source strength map then find a continuous trajectory under the influence of ocean currents.

Although the knowledge about the source is updated with concentration measurements over time, the knowledge is still uncertain. This leads to the *exploration-exploitation dilemma*. When there is not enough information about the source, the glider must *explore* the environment before it *exploits* the information.

In this section, we present a solution to the exploration-exploitation dilemma with *GP-upper confidence bound* (GP-UCB). The GP-UCB strategy adopts the principle of optimism under uncertainty, and picks the points that maximise the greatest value possible given the measurements so far (i.e. the upper confidence bound of posterior distribution). Note that the upper confidence bound is a weighted sum of predictive mean and standard deviation, which balances between exploitation (i.e., the mean) and the exploration (i.e., the variance).

In our case, we pick the next dive-to waypoint $\mathbf{x}_{k+1}$ as the one that maximises the following acquisition function:

$$\mathbf{x}_{k+1} = \underset{\mathbf{x} \in \mathcal{N}(\mathbf{x}_k)}{\arg\max}\, \mu_k(\mathbf{x}) + \beta_k \sigma_k(\mathbf{x}), \tag{24}$$

where $\mathcal{N}(\mathbf{x}_k)$ is a finite set of potential next dive-to locations, μ_k and σ_k are the predictive mean and standard deviation obtained from the GP after k samples, and β_k is a tuning parameter that balances between exploration and exploitation. As we will see later, β_k plays an important role in guaranteeing finding the source. Because we consider a finite set of potential next dive-to location, the acquisition function can be maximised by direct sampling.

Once the next best dive-to waypoint is found using GP-UCB, we find an energy-optimal glider trajectory in the form of a sequence of trim states as described in Sec. 3.1 using a sampling-based planning algorithm FMT*. The details about the sampling-based planning with trim state is previously presented in [Lee *et al.*, 2017].

The pseudocode for GP-UCB-based hierarchical planner is shown in Alg. 1. We first sample the candidate dive locations in line 5, and predict the source strength

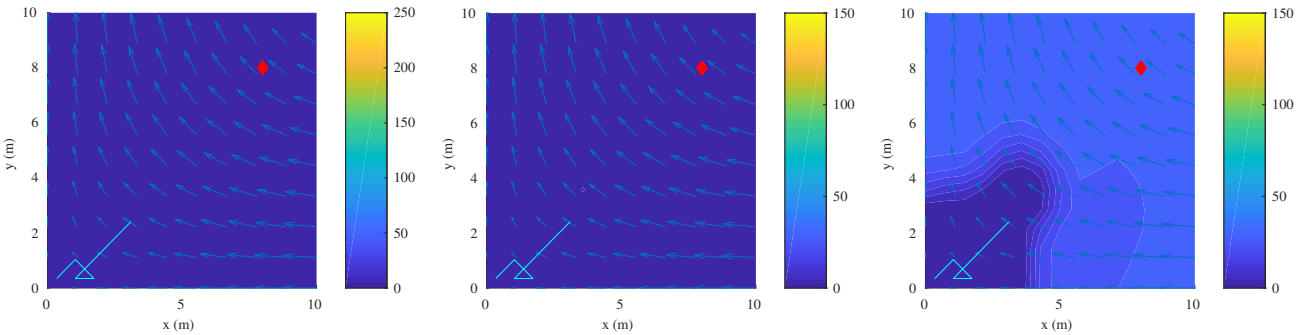

(a) $k = 6$, Est. concentration (b) $k = 6$, Est. source (c) $k = 6$, Source UCB

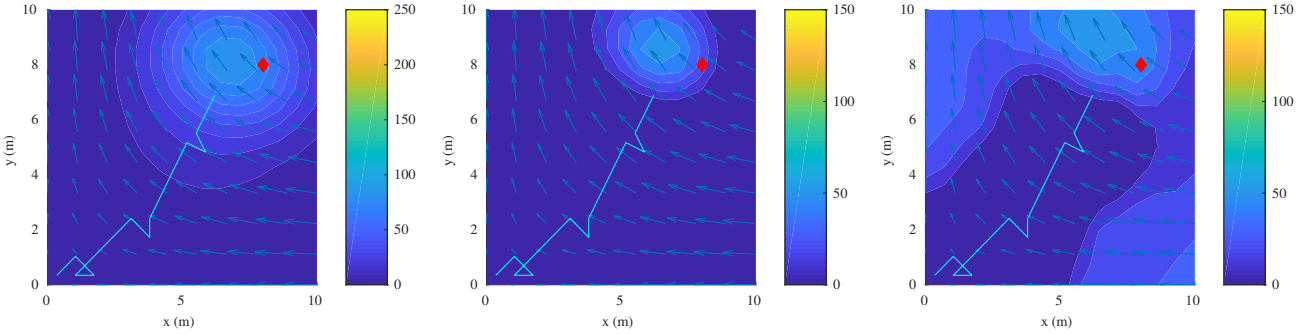

(d) $k = 16$, Est. concentration (e) $k = 16$, Est. source (f) $k = 16$, Source UCB

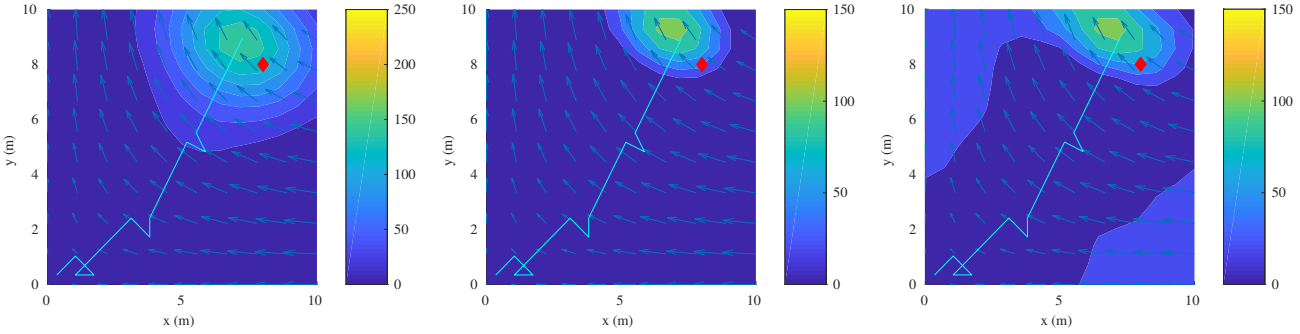

(g) $k = 19$, Est. concentration (h) $k = 19$, Est. source (i) $k = 19$, Source UCB

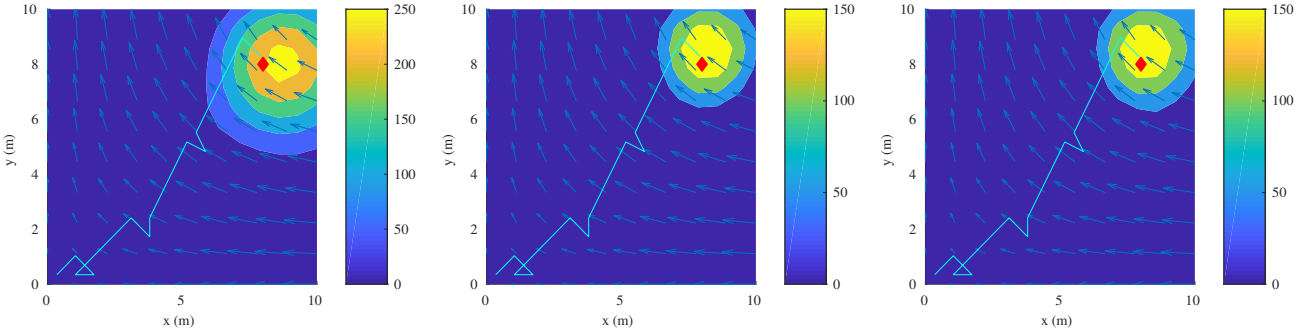

(j) $k = 20$, Est. concentration (k) $k = 20$, Est. source (l) $k = 20$, Source UCB

Figure 4: 2D simulation where glider (in light blue) manoeuvres through ocean currents (in pale blue) to find the source of plume (in red diamond)

Algorithm 1 GP-UCB-based Hierarchical Planner

1: GP ← InitialiseEmptyGP()
2: **while** Glider is operational **do**
3: **if** Previous dive $\mathbf{x}_k$ completed **then**
4: Update GP with measurements $\mathcal{O}_k$.
5: $\mathcal{N}_k \leftarrow$ GetNeighbours($\mathbf{x}_k$)
6: **for** $\mathbf{x} \in \mathcal{N}_k$ **do**
7: $\{\mu_s(\mathbf{x}), \sigma_s(\mathbf{x})^2\} \leftarrow$ GP.PredictSource($\mathbf{x}$)
8: **repeat**
9: $\mathbf{x}_{k+1} \leftarrow \arg\max_{\mathbf{x} \in \mathcal{N}_k} \mu_s(\mathbf{x}) + \beta_k \sigma_s(\mathbf{x})$
10: **until** $\mathbf{x}_{k+1}$ is feasible
11: $\mathbf{u}(t) \leftarrow \text{FMT}^*(\mathbf{x}_k, \mathbf{x}_{k+1})$.

at these candidate points (line 7). We pick the next dive location as the one that is feasible and maximises the GP-UCB acquisition function. After we find the best-next-waypoint $\mathbf{x}_{k+1}$, we find the glider controls from $\mathbf{x}_k$ using the FMT*algorithm.

6.1 Analysis

In an exploration-exploitation task where a decision is made over uncertain information, the decision may not be correct. The performance loss in making a wrong decision is referred to as *regret* [Srinivas *et al.*, 2010].

In the context of our source localisation task, regret is the disparity between dive-to location $\mathbf{x}_{k+1}$ and the true source location $\mathbf{x}_s$ in terms of source strength. In this paper, we define *average regret* R_k as the average of regrets over the sequence of 'dive-to location'-'true source location' pairs at k-th time step. Formally,

$$\bar{R}_k = \frac{1}{T} \sum_{i=1}^{k} s(\mathbf{x}_s) - s(\mathbf{x}_i). \tag{25}$$

The average regret approaches to zero asymptotically with high probability as path length increases. Intuitively, this is because the GP describes the correlation between observations and source strength, and ensures that the estimated source strength converges to the ground truth. In other words, using the source GP to estimate the source location becomes increasingly accurate. The convergence behaviour is satisfied with a particular choice of tuning parameter β_k. The following theorem provides the formal statement.

Theorem 1. *Pick* $\delta \in (0, 1)$*, and let*

$$\beta_k = 2\log(\frac{\pi^2 k^2 N_s}{6\delta}), \tag{26}$$

where N_s *is the size of potential sampling locations. Suppose we pick the sampling locations as per* (24)*. Then,*

$$\lim_{k \to \infty} \bar{R}_k = 0 \tag{27}$$

holds with probability greater than $1 - \delta$*.*

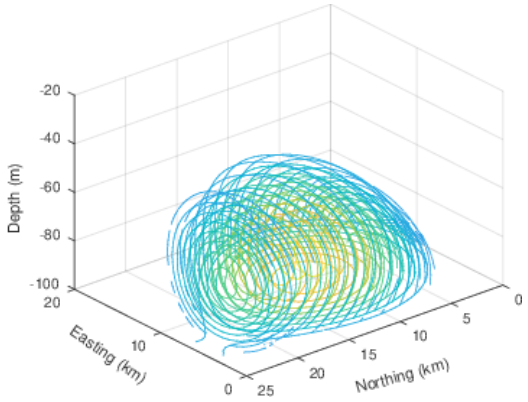

Figure 5: Ground truth concentration from experimental dataset shown with iso-value contours

Proof. The theorem is an application of Theorems 1 and 5 in [Srinivas *et al.*, 2010]. By Theorems 1 and 5 in [Srinivas *et al.*, 2010], we have $\bar{R}_k \leq \sqrt{\frac{C_1 \beta_k \gamma_T}{T}}$ for C_1 a constant with probability $> 1-\delta$, and $\gamma_T = O(\log(T)^3)$. In other words, $\exists k_0 \in \mathbb{N}$ and $\exists M_1 > 0$ such that $\forall k > k_0$, we have $\gamma_k < M \log(k)^3$. Thus, it holds that $\bar{R}_k \leq M_2 \sqrt{\frac{\beta_k \log k^3}{k}}$. Because $\sqrt{\frac{\beta_k \log k^3}{k}} \to 0$ as $k \to \infty$, the claim holds by the sandwich theorem. □

Intuitively, by Theorem 1, the GP-UCB-based planner guarantees that the true source location is eventually visited.

7 Case Studies

We demonstrate the case studies to show how our GP-UCB-based hierarchical planner performs with simulated environment and real dataset, where we used the dynamic parameters of Teledyne-Webb G2 Slocum glider [Webb *et al.*, 2001]. In both studies, the glider is initially given no prior knowledge over the true source location $\mathbf{x}_s$. We show that the glider actively balances between exploration and exploitation, and eventually reaches the source location.

7.1 Simulated Source Finding

In this section, we test the proposed GP-UCB algorithm on an simulated 2-dimensional environment with a double gyre-flow field and plume. The plume with respect to the flow field is generated using the advection-diffusion equation in (5).

The updated source strength estimations given a set of paths and observations are shown in Fig. 4. The glider initially has no prior knowledge of the source location shown in Fig. 4a. However, the UCB map allows the glider to choose an intelligent next dive-location, because the proposed source concentration GP can exploit the information about the flow field to compute a low

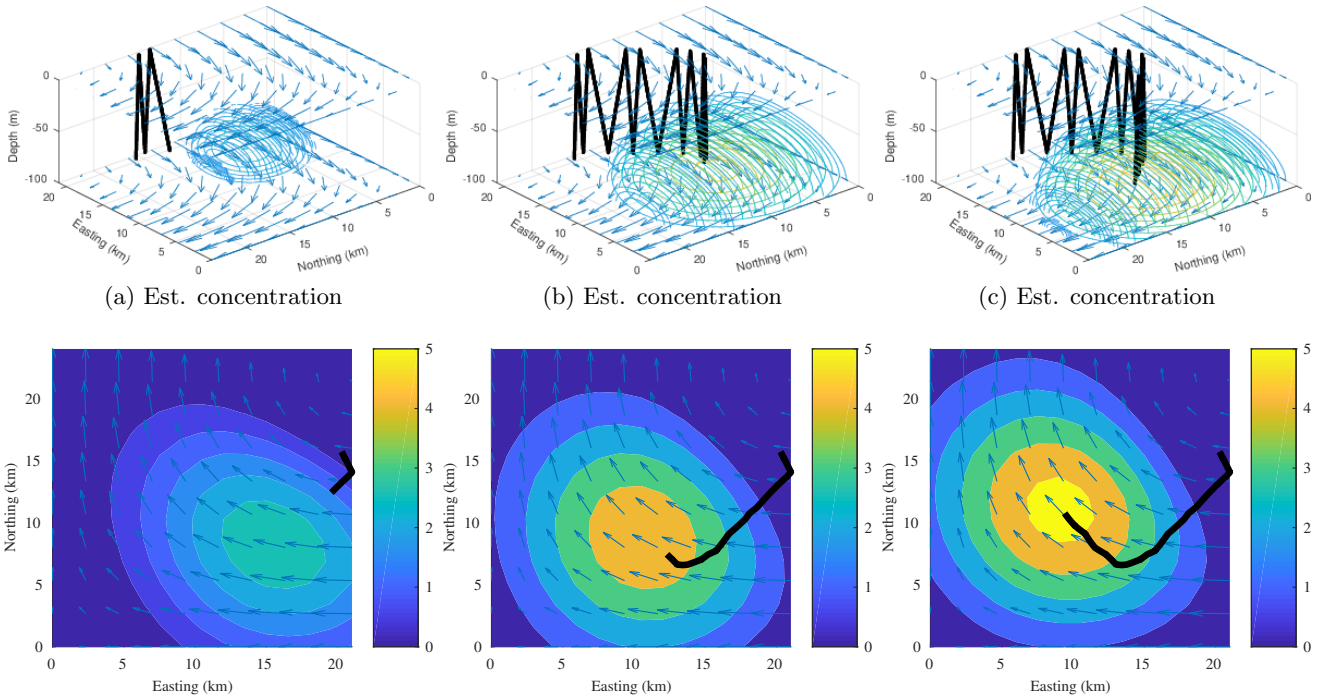

(a) Est. concentration (b) Est. concentration (c) Est. concentration

(d) Est. source (e) Est. source (f) Est. source

Figure 6: Planning over the real measurement dataset

uncertainty upstream of the glider. After a few exploration attempts, the glider gets close to a point (i.e., around $k = 16$), the GP-UCB started to estimate the belief over the concentration and the corresponding belief over source location, and show an interesting area to visit in the UCB map. As the glider moves further, the dive-in waypoint due to UCB approaches the true source location and the glider finally reaches the source.

An interesting observation is that when the upper confidence bound of the source strength becomes almost indistinguishable to the source strength when the estimate is converged. It is thus evident that the upper confidence bound correctly balances between exploration and exploitation. When the estimate is uncertain, the glider explores the areas where the uncertainty is higher. When the estimate is certain, the glider exploits the estimate and converges toward the true source location.

It is important to note that the estimate of the concentration does not necessary converge to the ground truth. This is because our objective is not to estimate the overall concentration of the overall environment, but to locate the source location. As such, the algorithm does not expend effort on trying to estimate the concentration correctly, unless it is necessary for estimating and converging toward the source location.

7.2 Real Measurement Dataset

We demonstrate the full 3D framework on an experimental dataset collected around Perth, Australia provided by Blue Ocean Monitoring [Russell-Cargill *et al.*, 2018] using G2 Slocum glider equipped with an optical methane gas sensor . The glider surveyed a region containing a known methane source identified previously. We pre-processed the collected data with GP regression to produce smooth 'ground truth' concentration, which is shown in Fig. 5. As the data is from the real ocean, it is unknown where the true source location is. However, it is intuitive that the glider expects to see the source at the centre of the plume.

The result of our proposed algorithm is shown in Fig. 5. It can be seen that the dive- and surface-locations are chosen sequentially to eventually reach the centre of the plume, by updating the concentration and source estimation. The actual true location of the source is not provided with the dataset. However, our analysis using GP-UCB shows that there is a high chance that the source is located around the centre of high UCB values. This is further ascertained by the convergence of the estimated source strength around the glider's final position.

8 Conclusion and Future Work

We have presented an algorithmic framework for localising the source of an underwater plume in an energy-efficient manner. The proposed algorithms are general enough to accommodate different types of flow fields and plume sources. In addition, the algorithms exhibit strong theoretical performance guarantee and good empirical performance. In the future, we would like to extend the framework to include further realistic constraints, such as energy or time budget, and to evaluate our work in the field.

Acknowledgments

This work is supported in part by an Australian Government Research Training Program (RTP) Scholarship, Blue Ocean Monitoring, the Australian Government through the Innovation Connections funding scheme (ICG000307),and the University of Technology Sydney.

References

[Cao *et al.*, 2017] Junliang Cao, Junjun Cao, Zheng Zeng, Baoheng Yao, and Lian Lian. Toward Optimal Rendezvous of Multiple Underwater Gliders: 3D Path Planning with Combined Sawtooth and Spiral Motion. *J. Intell. Robot. Syst.*, 85(1):189–206, 2017.

[Chang *et al.*, 2013] Dongsik Chang, Wencen Wu, Donald R Webster, Marc J Weissburg, and Fumin Zhang. A bio-inspired plume tracking algorithm for mobile sensing swarms in turbulent flow. In *Proc. of IEEE ICRA*, pages 921–926. IEEE, 2013.

[Fernández-Perdomo *et al.*, 2010] Enrique Fernández-Perdomo, Jorge Cabrera-Gámez, Daniel Hernández-Sosa, Josep Isern-González, Antonio C. Domínguez-Brito, Alex Redondo, Josep Coca, Antonio G. Ramos, Enrique Álvarez Fanjul, and Marcos García. Path planning for gliders using Regional Ocean Models: Application of Pinzón path planner with the ESEOAT model and the RU27 trans-Atlantic flight data. In *Proc. of IEEE OCEANS*, pages 1–10, 2010.

[Fernández-Perdomo *et al.*, 2011] Enrique Fernández-Perdomo, Daniel Hernández-Sosa, Josep Isern-González, Jorge Cabrera-Gámez, Antonio C. Domínguez-Brito, and Víctor Prieto-Marañón. Single and multiple glider path planning using an optimization-based approach. In *Proc. of IEEE OCEANS*, pages 1–10, 2011.

[Gunatilaka *et al.*, 2008] Ajith Gunatilaka, Branko Ristic, Alex Skvortsov, and Mark Morelande. Parameter estimation of a continuous chemical plume source. In *Proc. of IEEE ICIF*, pages 1–8. IEEE, 2008.

[Hajieghrary *et al.*, 2017] Hadi Hajieghrary, Daniel Mox, and M. Ani Hsieh. Information theoretic source seeking strategies for multiagent plume tracking in turbulent fields. *J. Mar. Sci. Eng.*, 5(1), 2017.

[Isern-González *et al.*, 2011] Josep Isern-González, Daniel Hernández-Sosa, Enrique Fernández-Perdomo, Jorge Cabrera-Gámez, Antonio C. Domínguez-Brito, and Víctor Prieto-Marañón. Path planning for underwater gliders using iterative optimization. In *Proc. of IEEE ICRA*, pages 1538–1543, 2011.

[Kowadlo and Russell, 2008] Gideon Kowadlo and R. Andrew Russell. Robot odor localization: A taxonomy and survey. *Int. J. Rob. Res.*, 27(8):869–894, 2008.

[Kularatne *et al.*, 2015] Dhanushka Kularatne, Ryan N. Smith, and M. Ani Hsieh. Zig-zag wanderer: Towards adaptive tracking of time-varying coherent structures in the ocean. *Proc. of IEEE ICRA*, 2015-June(June):3253–3258, 2015.

[Lee *et al.*, 2017] James Ju Heon Lee, Chanyeol Yoo, Raewyn Hall, Stuart Anstee, and Robert Fitch. Energy-optimal kinodynamic planning for underwater gliders in flow fields. *Proc. of ACRA*, pages 42–51, 2017.

[Leonard and Graver, 2001] N E Leonard and J G Graver. Model-based feedback control of autonomous underwater gliders. *IEEE J. Ocean. Eng.*, 26(4):633–645, 2001.

[Li *et al.*, 2014] S. Li, Y. Guo, and B. Bingham. Multi-robot cooperative control for monitoring and tracking dynamic plumes. In *Proc. of IEEE ICRA*, pages 67–73, May 2014.

[Liu *et al.*, 2017] Yanji Liu, Jie Ma, Ning Ma, and Guichen Zhang. Path planning for underwater glider under control constraint. *Adv. Mech. Eng.*, 9(8):1–9, 2017.

[Lupini and Tirabassi, 1979] R. Lupini and T. Tirabassi. Gaussian plume model and advection-diffusion equation: An attempt to connect the two approaches. *Atmos. Environ.*, 13(8):1169 – 1174, 1979.

[Marthaler and Bertozzi, 2004] Daniel Marthaler and Andrea L Bertozzi. Tracking environmental level sets with autonomous vehicles. In *Recent developments in cooperative control and optimization*, pages 317–332. Springer, 2004.

[Mrazovac *et al.*, 2012] Sanja M. Mrazovac, Pantic R. Milan, Mirjana B. Vojinovic-Miloradov, and Bratislav S. Tosic. Dynamic model of methanewater diffusion. *Appl. Math. Model.*, 36(9):3985 – 3991, 2012.

[Rao and Williams, 2009] Dushyant Rao and Stefan B. Williams. Large-scale path planning for Underwater Gliders in ocean currents. In *Proc. of ACRA*, 2009.

[Rasmussen and Williams, 2006] Carl Edward Rasmussen and Christopher K. I. Williams. *Gaussian processes for machine learning*. The MIT Press, 2006.

[Rudnick *et al.*, 2004] Daniel L Rudnick, Russ E Davis, Charles C Eriksen, David M Fratantoni, and Mary Jane Perry. Underwater gliders for ocean research. *Mar. Technol. Soc. J.*, 38(2):73–84, 2004.

[Russell-Cargill *et al.*, 2018] Louise M. Russell-Cargill, Bradley S. Craddock, Ross B. Dinsdale, Jacqueline G. Doran, Ben N. Hunt, and Ben Hollings. Using autonomous underwater gliders for geochemical exploration surveys. *The APPEA Journal*, 58:367–380, 2018.

[Särkkä and Hartikainen, 2012] Simo Särkkä and Jouni Hartikainen. Infinite-dimensional kalman filtering approach to spatio-temporal gaussian process regression. In *Proc. of AISTATS*, pages 993–1001, 2012.

[Shih *et al.*, 2017] Chien-Chou Shih, Mong-Fong Horng, Tien-Szu Pan, Jeng-Shyang Pan, and Chun-Yu Chen. A genetic-based effective approach to path-planning of autonomous underwater glider with upstream-current avoidance in variable oceans. *Soft Comput.*, 21(18):5369–5386, 2017.

[Smith *et al.*, 2010] Ryan Smith, Yi Chao, Peggy Li, David A. Caron, Burton Jones, and Gaurav Sukhatme. Planning and implementing trajectories for autonomous underwater vehicles to track evolving ocean processes based on predictions from a regional ocean model. *Int. J. Rob. Res.*, 29:1475–1497, 10 2010.

[Srinivas *et al.*, 2010] Niranjan Srinivas, Andreas Krause, Sham Kakade, and Matthias Seeger. Gaussian process optimization in the bandit setting: No regret and experimental design. In *Proc. of ICML*, ICML'10, pages 1015–1022, USA, 2010. Omnipress.

[Vergassola *et al.*, 2007] Massimo Vergassola, Emmanuel Villermaux, and Boris I Shraiman. infotaxis' as a strategy for searching without gradients. *Nature*, 445:406, jan 2007.

[Webb *et al.*, 2001] Douglas C. Webb, Paul J. Simonetti, and Clayton P. Jones. SLOCUM: An underwater glider propelled by environmental energy. *IEEE J. Ocean. Eng.*, 26(4):447–452, 2001.

[Xinke *et al.*, 2015] Zhu Xinke, Jin Xianglong, Yu Jiancheng, and Li Yiping. Path planning in stronger ocean current for underwater glider. In *Proc. of IEEE CYBER*, pages 891–895, 2015.

[Zamuda and Hernández Sosa, 2014] Aleš Zamuda and José Daniel Hernández Sosa. Differential evolution and underwater glider path planning applied to the short-term opportunistic sampling of dynamic mesoscale ocean structures. *J. Appl. Soft Comput.*, 24:95–108, 2014.