Evaluation of AUV Search Strategies for the Localization of Hydrothermal Venting

[1]James C. Mason, [1]Andrew Branch, [2]Guangyu Xu, [3]Michael V. Jakuba,
[3]Christopher R. German, [1]Steve Chien, [3]Andrew D. Bowen, [1]Kevin P. Hand, [3]Jeffrey S. Seewald
[1]Jet Propulsion Laboratory, California Institute of Technology
[2]Applied Physics Laboratory, University of Washington
[3]Woods Hole Oceanographic Institution
Corresponding Author: andrew.branch@jpl.nasa.gov

Abstract

Ocean Worlds represent one of the best chances for the discovery of extra-terrestrial life within our own solar system, particularly near sources of hydrothermal venting. To study the oceans on Ocean Worlds will require a new type of mission to penetrate the icy shell, deploy an autonomous underwater vehicle (AUV), and travel potentially hundreds of kilometers with minimal contact to Earth based operations teams. To maximize the science return, the AUV would need to be capable of fully autonomously locating and studying scientific features of interest. We have developed two strategies to locate sources of hydrothermal venting: a gradient ascent strategy and a greedy transect search strategy. We have improved a previously-implemented nested search strategy by adding a vertical search component. Each strategy is tested in a hydrothermal plume dispersion simulation. We compare the effectiveness of each method in this environment.

Introduction

At least eight bodies in our solar system are thought to harbor liquid oceans (National Aeronautics and Space Administration 2018). All of which, with the exception of Earth, are covered in an icy shell. To explore these worlds an ice penetrating submersible vehicle is required. A notional mission concept for such a submersible (Branch et al. 2018), contains three main components: a surface antenna, an under-ice base station, and a submersible vehicle. When the surface antenna does not have line of sight with Earth, 50% of the orbital period or 42 out of every 85 hours in the case of Europa , or when the vehicle is out of acoustic communication range with the under-ice base station, which will be necessary to explore a large portion of the ocean, there will be no communication with Earth. To maximize the science return of the mission the submersible would be required to autonomously detect, locate, and study a specific feature of interest (Chien and Wagstaff 2017).

Hydrothermal venting is one potential scientific target for a submersible mission. Evidence for hydrothermal activity has been found on Enceladus (Hsu et al. 2015; Waite et al. 2017). On Earth, hydrothermal vents harbor unique ecosystems and are potentially critical to the origin of life (Jordan et al. 2019). Similar phenomena on Ocean Worlds could be the best chance at extra-terrestrial life in our Solar System. Europa also shows great potential for habitability (Hand et al. 2009; Trumbo, Brown, and Hand 2019) and may also host plume activity directly comparable to that already observed on Enceladus (Sparks et al. 2016; 2017; Jia et al. 2018).

We have newly implemented two autonomous strategies for the localization of hydrothermal vents. A horizontal gradient ascent strategy and a greedy transect search strategy, both inspired by (Burian et al. 1996). We have improved a previously-developed nested search strategy (Branch et al. 2018) by including a vertical search component. The vertical search is also used in the greedy transect strategy. These approaches are compared in a simulation environment developed in (Branch et al. 2018). As in (Branch et al. 2018), we focus on search in the non-buoyant hydrothermal plume.

The rest of the paper is organized as follows. First we discuss related work and describe the structure of hydrothermal venting. Then we outline the various search strategies. We describe the simulation environment used to test our approach and the experimental setup. Finally we discuss the results and future work.

Related Work

Our work directly extends from (Branch et al. 2020) and (Branch et al. 2018) in which we previously tested a single method, the nested search strategy. We now present two new strategies. We previously did not perform vertical search outside of the spiral phase of the nested search – we now perform vertical search through the entire nested bin strategy.

Foundational work in adaptive sampling and control of autonomous underwater vehicles was done by the Autonomous Ocean Sampling Network (Curtin et al. 1993; Curtin and Bellingham 2009; Ramp et al. 2009; Haley et al. 2009; Leonard et al. 2007).

A non-autonomous three-phase nested search is currently done on Earth to locate sources of hydrothermal venting (German et al. 2008; Yoerger et al. 2007a). This has been augmented by (Yoerger et al. 2007b) to autonomously revisit areas of interest after a primary mission has been completed, however this required humans to design the primary mission. This was used in the field multiple times. (Farrell, Pang, and Li 2005) developed a method to follow chemical plumes via a strategy inspired by moths. A number of other

Copyright © 2020, All rights reserved.

hydrothermal vent localization strategies have been implemented and tested in idealized simulation environments or in post-processing with deployment data. (Pang 2010) and (Tian et al. 2014) use moth based strategies, (Jakuba and Yoerger 2008) uses occupancy grid, (Saigol et al. 2010) uses a belief-maximization algorithm, and (Ferri, Jakuba, and Yoerger 2010) uses a trigger based approach to perform higher resolution surveys in regions of interest.

Gradient following has been used with AUVs to identify the global maximum in some measure; in (Burian et al. 1996) they locate the maximum depth and maximum temperature in a pond with a search area of 125m by 143m with a theoretically maximum search duration of 68 minutes(Burian et al. 1996). Our search area covers approximately 60km by 60km and terminates if it is still searching after 28 days of simulated time. The plume strength at a particular (x, y, z) location may change dramatically during this period, so we are performing gradient ascent on a constantly changing function. (Burian et al. 1996) also developed the greedy transect strategy and demonstrated it in the pond as described with the gradient strategy. As stated earlier, our work is a simulation that covers an extremely large search area, while their work was a technology demonstration in a small pond.

Hydrothermal venting is not the only target of interest. While not all ocean processes on Earth are expected to recur on other ocean worlds distant from the sun, we have a wealth of experience studying thermoclines, ocean fronts, and other structures in Earth's oceans. In targeting these specific scientific features the objective can generally be categorized as locating the point or boundary between two regions of water with the highest gradient. A number of different near real-time feature tracking methods exist for thermoclines (Cruz and Matos 2010; Zhang et al. 2010; Sun et al. 2016). (Zhang et al. 2013; 2016) tracks upwelling fronts using a zig-zag pattern. (Cruz and Matos 2014) tracks any gradient boundary using a single vehicle following a dynamic zig-zag pattern and a lateral gradient detection algorithm to estimate the gradient boundary using an arc. A similar method can also be applied to tracking the center of a phytoplankton bloom patch (Godin et al. 2011). (Branch et al. 2019) uses near real-time data to autonomously retask a set of vehicles to repeatedly sample an ocean front. The goal of this work is to maximize the crossing of a front. This differs from our work which focuses on finding the maximum of a single point source. Other work focuses on tracking algal blooms by flying formations relative to the bloom as tracked by a drifter (Das et al. 2012). (Petillo, Schmidt, and Balasuriya 2012) uses a simulated network of AUVs in order to estimate the boundary of a simulated plume. (Flexas et al. 2018) uses an ocean model and autonomous planning to optimize sampling of submesoscale structures.

Hydrothermal Venting

Seafloor hydrothermal vents produce a plume of chemically altered seawater that can be used to locate the source. The plume structure is shown in Figure 1. The less dense hydrothermal fluid exiting the vent forms the buoyant plume. As this plume rises it is continuously diluted by the ambient seawater, expanding from approximately 10 cm at the vent to about 100 meters when it reaches equilibrium. At equilibrium the plume extends horizontally, on the order of 10s of kilometers, to form the non-buoyant plume (German and Seyfried 2014). The height of the non-buoyant plume is variable as it is a function of the vent fluid and the surrounding seawater (Turner 1979), however it is on the order of 100s of meters (Speer and Rona 1989).

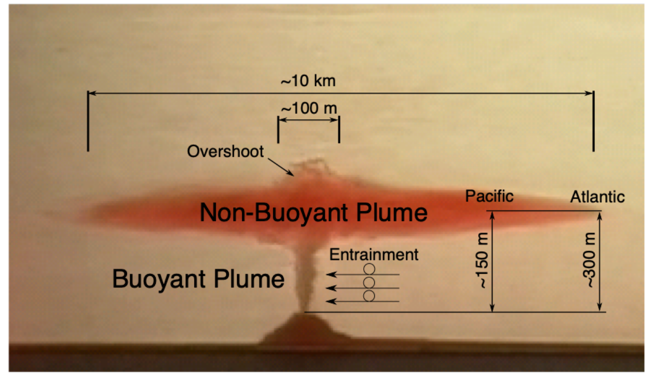

Figure 1: Demonstration of a hydrothermal plume performed in an aquarium tank. The buoyant and non-buoyant components of the hydrothermal vent plume are labeled with approximate scales. Image courtesy of C. German, WHOI

The hydrothermal plume is the primary source of information for the vehicle to localize the venting source. This is complicated by tidal flows, causing local maxima unassociated with the vent source (Veirs 2003), turbulent flow preventing smooth gradients from forming, and unknown vent properties. Chemical sensors, such as oxidation-reduction potential, are the primary sensors used to locate hydrothermal vent sources (Nakamura et al. 2000), augmented by temperature and optical backscatter (Baker, German, and Elderfield 1995; Baker and German 2004). These sensors may be good candidates for inclusion on a submersible mission to an Ocean World due to their compact form factor (100s of grams) and low power consumption (10s of milliwatts).

Search Strategies

Given a vehicle's starting location, the goal is to sample to the hydrothermal plume and produce a control strategy that results in locating the vent source. Performance is measured based on the horizontal distance from the vehicle's predicted source location to the true source location. For this paper a successful run will be within 200 meters, limited by the resolution of the hydrothermal venting model used. Performance is also based on the time required to complete the search. While hydrothermal plumes (and our model) vary temporally, we assume that this variation is not significant to consider while searching.

We implemented three strategies: the Gradient ascent strategy, the Greedy Transect strategy, and the Nested Search strategy.

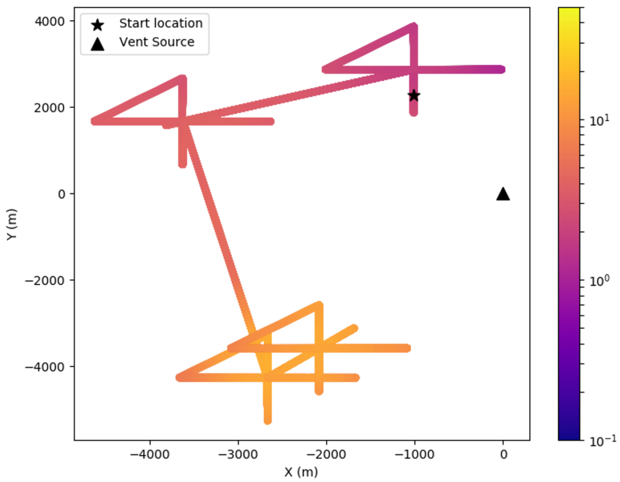

Figure 2: Top down plot displaying the passive tracer as seen in a scenario for the gradient ascent strategy. The vehicle performs transects in the x and y directions, then follows the resulting gradient until it becomes negative. We stop searching when maximum observed tracer is close to our previous point. The spiral phase is omitted for clarity.

Gradient Ascent Search Strategy

When searching for maxima in complex functions, the first approach is often a gradient Ascent strategy. We implemented a gradient ascent strategy, inspired by (Burian et al. 1996), that samples the passive tracer in the x and y directions, and combines the results in a gradient vector. The approach is outlined in Algorithm 1 and operates as follows. A spiral is initiated at the start location, completing vertical yoyos as it travels. When the max plume strength value of a single profile exceeds the specified threshold, we then begin the gradient ascent search. At a constant depth, the AUV follows two transects of length $2 * transect\_length$, in the x direction then the y direction, and computes two gradients, $grad_x$ and $grad_y$. Then atan2 is used to compute a single 2d gradient, $grad\_dir$. The gradient is followed, still at a constant depth, for some minimum distance and until the gradient, calculated via the slope of a least squares linear fit on a sliding window, becomes negative. If the location of the maximum observed tracer along the gradient following is within $final\_spacing$ of the previous start location, we stop searching.

Greedy Transect Search Strategy

We also implemented a greedy transect strategy (Burian et al. 1996). The approach is outlined in Algorithm 2 and operates as follows. A spiral is initiated at the start location, completing vertical yoyos as it travels. When the max plume strength value of a single profile exceeds the plume threshold, we then begin performing transects in a star pattern, called $star\_transects$. The transects are centered on the location of the maximum observed value so far, (x_d, y_d). A total of 8 transects are performed, all starting at the center

Algorithm 1 Gradient Ascent Search

```
procedure GRADIENT_SEARCH
  plan ← spiral
  while plan ≠ NULL and not timed out do
    Execute or Continue plan
    if executing spiral then
      Wait until end of vertical profile
      p_d ← Get data from profile
      d ← max(p_d)
      if d ≥ plume threshold then
        (x_d, y_d) ← location(d)
        plan ← transect_x at x_d between y_d ± transect_length
    else if executing transect_x then
      Wait until end of the transect
      p_d ← Get data from transect
      grad_x ← slope of linear least squares fit of p_d
      plan ← transect_y at y_d between x_d ± transect_length
    else if executing transect_y then
      Wait until end of the transect
      p_d ← Get data from transect
      grad_y ← slope of linear least squares fit of p_d
      grad_dir ← atan2(grad_x, grad_y)
      plan ← gradient_follow from (x, y) in direction grad_dir
    else if executing gradient_follow then
      wait until slope of least squares linear fit of line data < 0 and traveled min distance
      p_d ← Get data from profile
      d ← max(p_d)
      (x_f, y_f) ←location(d)
      if dist(location(d), (x_d, y_d)) ≤ final_spacing then
        return Success
      (x_d, y_d) ← (x_f, y_f)
      plan ← transect_x at x_d between y_d ± transect_length
  return Failure
```

with orientations spaced at 45 degree intervals. Transects are always performed in pairs, where a pair is a transect and its opposing transect (e.g North and South or East and West). After completing a pair of transects and if we observed a new maximum, we repeat the transects at the location of that new maximum. For a given transect, the vehicle travels along that transect until the plume strength has been observed to be decreasing. This is determined by binning the data along the transect, taking the average of each bin, and completing the transect if the last $num\_section\_threshold$ bins have been decreasing. These calculations are performed in the SHOULD_END_TRANSECT procedure. If we have completed all transects without observing a larger maximum then the search terminates. The center point of the last pattern is returned as our estimated vent location. The AUV performs vertical search (see Algorithm 5) for the entire duration of the algorithm.

Nested Search Strategy

The description of the nested search algorithm from (Branch et al. 2020) and (Branch et al. 2018) is included here and outlined in Algorithm 3. To initially locate the plume a spiral pattern with vertical yoyos is performed centered on the start location. At the end of each profile, if the max plume strength value seen on that profile exceeds the threshold, $plume_t$, we move to the next portion of the search process.

The search space is divided into four quadrants centered on the location where the plume was detected in the previous step. These quadrants are then further partitioned into bins, $survey\_bins$ in Algorithm 3, of size $spacing_0$. A dynamic "lawnmower" survey is executed in each of the four quadrants. The dynamic lawnmower survey dynamically resizes based on the observed data. It is outlined in Algorithm 4. The spacing of the lawnmower pattern, $track\_spacing$ is specified beforehand, based on the expected size of the

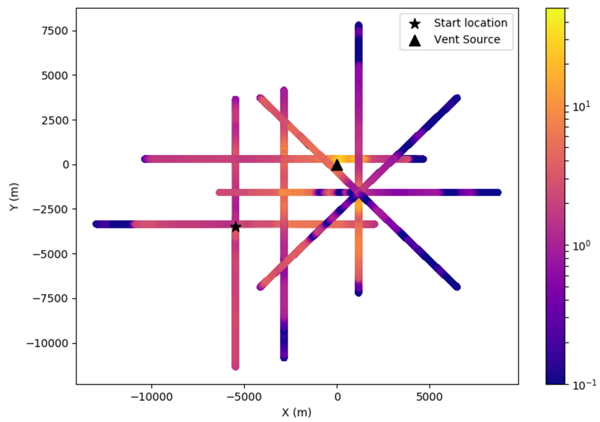

Figure 3: Top down plot displaying the passive tracer as seen in a scenario for the greedy transect strategy. The vehicle performs transects in a star pattern, restarting the pattern whenever it discovers a new maximum. We stop searching when no new maxima are discovered. There is some variation in the data for a transect because we are performing yoyos, causing us to sample at various depths. The spiral phase is omitted for clarity.

Algorithm 2 Greedy Transect Search

```
procedure TRANSECT_SEARCH
    plan ← spiral
    while plan ≠ NULL and not timed out do
        Execute or Continue plan
        if executing spiral then
            Wait until end of vertical profile
            p_d ← Get data from profile
            d ← max(p_d)
            if d ≥ plume threshold then
                (x_d, y_d) ← location(d)
                plan ← star_transects centered at (x_d, y_d)
        else if executing star_transects then
            Wait until SHOULD_END_TRANSECT(transect data)
            if A pair of transects is complete then
                if there is a new maximum in the transect then
                    (x_d, y_d) ← location(new max)
                    plan ← star_transects centered at (x_d, y_d)
                else if we have completed star_transects then
                    return Success
                else
                    continue
            else
                continue
    return Failure
procedure SHOULD_END_TRANSECT(data)
    if have not traveled min_leg_dist then
        return False
    else if have traveled max_leg_dis then
        return True
    for each section i of size leg_section_length do
        section_avg[i] = average of section i
    for the last num_section_threshold − 1 sections do
        if section_avg[i] > section_avg[i − 1] then
            return False
    return True
```

features of interest. $along\_track$ and $across\_track$ define the two axis on which the lawnmower pattern will be performed. The lawnmower track lines are partitioned into sections of length equal to the spacing of the lawnmower. At least $min\_sections$ sections are be completed per track line. The plume strength values in each bin are averaged. If the average for the last $sections\_limit$ sections are below $plume_t$ and are monotonically decreasing, then the track line is completed and the vehicle moves to the next track line. $min\_sections$ and $sections\_limit$ are manually specified search parameters. If the maximum value of an entire track line is less than $plume_t$ then the current lawnmower survey is ended and the next begins. The data collected during the dynamic lawnmowers is binned into $survey\_bins$ for the next step.

Once a dynamic lawnmower has been completed, we check for any local maxima in $survey\_bins$. A local maximum is defined as when the 8 neighboring bins of the same partition size have a max plume detection less that that of the center bin. Nested lawnmower surveys are then performed at any maxima that has been found. The local maximum bin and all 8 neighboring bins are partitioned into thirds. A lawnmower pattern with spacing equal to these new smaller partitioned bins is then performed on those bins. This pattern will also cover the neighbors to the local maximum. If we have found multiple maximum then they are prioritized based on plume strength, however all maximum will be investigated before continuing with the next dynamic lawnmower pattern. This process repeats recursively, with nested lawnmowers of smaller resolution, until a resolution of $final\_spacing$ is reached. In real world operations the spiral pattern would resume after all dynamic lawnmowers have been completed. However, in our case, with a simulated hydrothermal vent source, the search terminates (Branch et al. 2020; 2018).

Yoyo Vertical Search

The algorithm descriptions thus far have focused on search in the X-Y (horizontal) plane. Search is also performed in the Z (vertical) axis. As with the other search methods we assume the sampled data is not time varying. The height of the plume varies depending on the properties of the water column at any given location. By searching in the Z direction we can maintain contact with the strongest part of the plume. Searching the entire ocean depth isn't feasible or desired, we want to stay close to the center of the plume. To do so we perform dynamic vertical yoyos, described in algorithm 5.

During the search the vehicle will oscillate through the water column while calculating the gradient of the plume strength in a sliding window, p_d, of size $grad\_time$. The gradient is calculated by taking the slope of the linear least squares fit of the data. If the gradient becomes negative, the vehicle will turn around in the vertical direction. When the AUV isn't in the plume, there may be no gradient to follow. In this case, the vehicle searches between the shallowest and deepest depths where it observed a tracer stronger than the $plume\_threshold$ parameter.

Simulation

We used a previously developed simulation environment that uses a hydrothermal plume dispersion simulation and a vehicle model (Branch et al. 2018). A numerical simulation of hydrothermal plume dispersion is performed using FVCOM, an ocean-circulation model, at Axial Seamount on the Juan de Fuca Ridge. FVCOM is a finite-volume, time

Algorithm 3 Nested Search

```
procedure NESTED_SEARCH
    plans ← empty stack
    visited ← empty set
    plans.push(spiral)
    survey_bins ← bins of size spacing_0
    while plans.size > 0 and not timed out do
        Execute or Continue plans.top()
        if executing spiral then
            Wait until end of vertical profile
            p_d ← Get data from profile
            d ← max(p_d)
            if d >= plume_t and d.location not explored then
                bins ← profile_data binned at 10 meters and averaged
                p_h ← max(bins).height
                (x, y) ← bin corner closest to d.position
                plans.push(dynamic_lawnmower(x, y, p_h, 90°, 0°, spacing_0))
                plans.push(dynamic_lawnmower(x, y, p_h, −90°, 0°, spacing_0))
                plans.push(dynamic_lawnmower(x, y, p_h, −90°, 180°, spacing_0))
                plans.push(dynamic_lawnmower(x, y, p_h, 90°, 180°, spacing_0))
                Execute plans.top()
        else
            while plans.top() is not completed do
                Wait
            survey_data ← Get data from latest survey
            survey_bins.add_data(survey_data)
            maxima ← get_bin_maxima(survey_bins)
            sort maxima
            for bin in maxima do
                if bin not in visited then
                    Partition bin and bin.neighbors()
                    visited.add(bin)
                    plans.push(nested_lawnmower(bin))
                    break
        while plans.size > 0 and plans.top() is complete do
            f ← plans.pop()
            if f.spacing < final_spacing and f contains vent source then
                return Success
    return Failure
```

Algorithm 4 Dynamic Lawnmower

```
procedure DYNAMIC_LAWNMOWER(x, y, h , along_track, across_track, track_spacing)
    start_x ← x + cos(along_track) ∗ track_spacing/2
    start_y ← y + sin(across_track) ∗ track_spacing/2
    Go to (start_x, start_y, h)
    curr_track ← 0
    curr_section ← 0
    completed ← False
    section_data ← empty list
    Start current track line on heading along_track
    while not completed do
        Do next section on current track
        section_data[curr_section] ← Get data from last section
        curr_section ← curr_section + 1
        if curr_section >= min_sections or survey boundary reached then
            if avg(section_data[i]) < plume_t for last sections_limit sections and
monotonically decreasing then
                curr_track ← curr_track + 1
                if max(section_data) < plume_thresh then
                    completed ← True
                section_data ← empty list
                Travel track_spacing on heading across_track
                if curr_track is even then
                    Start next track line on heading along_track
                else
                    Start next track line on heading −along_track
```

Algorithm 5 Yoyo Vertical Search

```
procedure YOYO_VERTICAL_SEARCH
    plan ← go_down
    while plan ≠ NULL and not done do
        Execute or Continue plan
        p_d ← Get data collected in the last grad_time seconds
        if  slope of linear least squares fit of p_d < 0 then
            if executing go_down then
                plan ← go_up
            else
                plan ← go_down
```

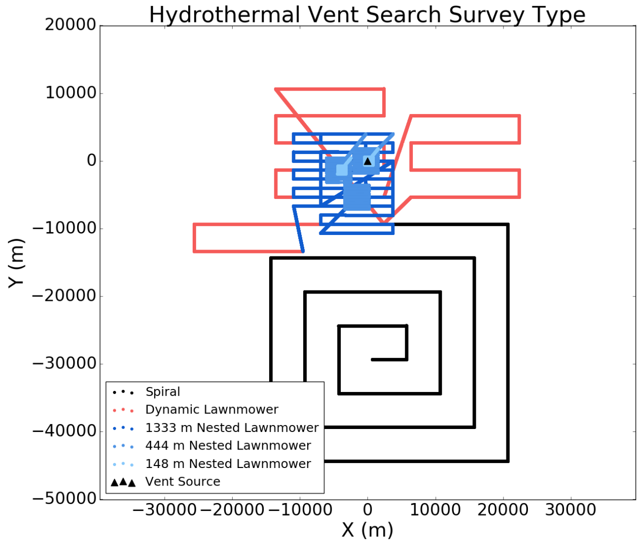

Figure 4: Top down plot displaying the types of surveys performed for the nested bin strategy. The search space is divided into square bins. Upon finding a local maximum, a lawnmower pattern is executed on a row of bins. The search starts near $(15000, -15000)$. The vent source is at $(0, 0)$.

and density-dependent, three-dimensional, ocean circulation model (Chen, Liu, and Beardsley 2003). In addition, FVCOM supports the use of large-scale ocean circulation and tidal model outputs as open boundary forcing to drive flow across a broad range of frequencies inside the model domain (Zheng and Weisberg 2012). The model domain is 300 by 300 km and centered on the Axial Seamount caldera. The horizontal resolution is variable, from 200 meters in a 10 by 10 km region over Axial's caldera to 10km at the domain boundary. The vertical component has 127 uniformly distributed layers. This results in a variable vertical resolution, with about 12 meter layer thickness above Axial's caldera. The duration of the simulation is 58 days, with model outputs sampled hourly. The model initializes stratification and performs open boundary forcing via the HYbrid Coordinate Ocean Model (HYCOM) and OSU Tidal Inversions models. Surface wind forcing and heat flux are incorporated from 1-hourly sampled National Centers for Environmental Prediction (NCEP) Climate Forecast System Reanalysis (CFSR) outputs. A seafloor heat source of 1 GW, representing the hydrothermal vent, in placed at the center (0,0) of the model domain. The model output consists of current, temperature, salinity, and a passive tracer, which is released at the vent source. In our experiments the passive tracer is used as a proxy for plume strength, with a value range of $[0, 100]$. The tracer reaches a quasi-steady state after 30 days of model simulation in a 20 by 20 km region centered on the vent source. No quasi-steady state is reached in a 50 by 50 km region before the end of the simulation. Therefore we begin our simulations at the models 30 day mark.

The nominal vehicle speed is set to 0.5 m/s. Simulated sensors are used to measure temperature, the passive tracer,

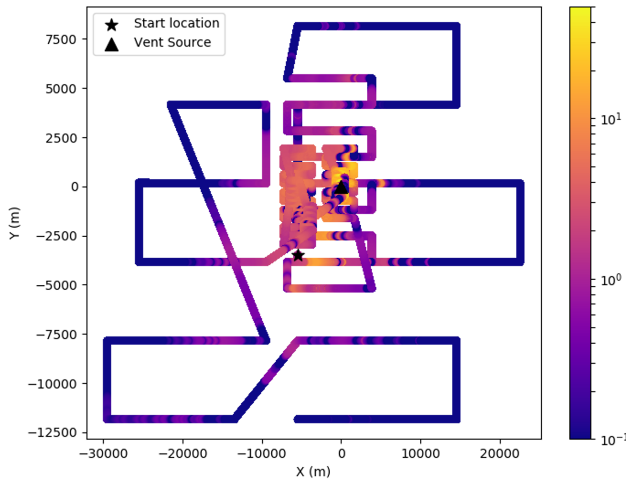

Figure 5: Top down plot displaying the passive tracer as seen in a scenario for the nested bin strategy. There is some variation in the data for a transect because we are performing yoyos, causing us to sample at various depths. The spiral phase is omitted for clarity.

vehicle depth, and distance to seafloor at a fixed interval. The position of the vehicle is assumed to be known at all times. Currently a chemical sensor, such as oxidation-reduction potential, and vehicle resources, such as energy and data capacity, are not modeled (Branch et al. 2018).

Experiment

25 scenarios were completed with the vehicle starting location uniformly varied with $x \in [-30000\,\text{m}, 30000\,\text{m}]$ and $y \in [-30000\,\text{m}, 30000\,\text{m}]$ at intervals of 15000 m. The domain of the simulation is $x \in [-150000\,\text{m}, 150000\,\text{m}]$ and $y \in [-150000\,\text{m}, 150000\,\text{m}]$. Due to the nature of some of the algorithms and the location of the vent at $(0, 0)$ it is likely that the vehicle will pass directly over the vent source if the start location x and y are multiples of 1000. To mitigate this, a uniformly random value between $[-1500, 1500]$ was added to the x and y values of the starting location. The simulated vehicle has a horizontal velocity of 0.5 m/s, and a vertical velocity of 0.25 m/s. The vehicle samples the model at 0.2 hz.

The plume detection threshold was set to 0.5. The initial spiral spacing was set to 5000 m and the initial dynamic lawnmower spacing was set to 4000 m. The dynamic lawnmower parameters $min\_sections$ and $sections\_limit$ are set to 4 and 2, respectively.

Results

We focus on two statistics: the distance between the true vent location and the predicted vent location at the end of a run (see Table 6), and how long in simulated time it took to complete the search (see Table 7). Figure 8 compares these two statistics. A run times out if it takes longer than 28 days.

Strategy	Min	Q1	Q2	Q3	Max
Nested	11.7	56.6	93.8	151.5	11520.8
Transect	34.8	70.1	128.1	1724.0	17383.3
Gradient	61.5	2915	6083.5	7415.0	20068.1

Figure 6: Five number summary for the distance (m) between the real vent location to the predicted vent location at the end of a run. We want to be within a few hundred meters. Q1, Q2, and Q3 describe the values which one quarter, two quarters, and three quarters of the data fall below.

Strategy	Min	Q1	Q2	Q3	Max
Nested	9.7	24.8	27.2	28.0	28.0
Nested Max	6.7	11.5	16.8	21.4	23.6
Transect	3.0	6.9	14.2	19.5	24.6
Gradient	2.0	3.2	8.3	13.2	18.9

Figure 7: The runtime (days) that the vehicle searched for, including runs where the vehicle didn't find the source or timed out. The search times out after 28 days. Nested Max is the time that the AUV observed the max plume strength for the Nested search strategy. Q1, Q2, and Q3 describe the values which one quarter, two quarters, and three quarters of the data fall below.

The nested search strategy's predicted source location was most frequently close to the true location. 80% of the 25 problems returned a predicted plume source location within 200m of the source. The nested search took the longest to complete of all the algorithms, with 75% of the runs taking longer than 24 days. However, the AUV frequently finds the maximum several days before terminating. This suggests that we may need to improve the strategy's stopping condition or more selectively investigate maxima.

The greedy transect strategy performed well at finding the source, predicting a vent within 200m of the source 56% of the time. The greedy transect strategy finished relatively quickly, with 75% of the runs taking less than 20 days.

The gradient ascent strategy performed poorly, with only 1 out of 25 runs finishing within 200m of the source. It was over 3km away from the source for 72% of the runs. When looking at the data, we noticed that it typically found a local maximum but was unable to find the global maximum. This is unsurprising as the gradient around the local maxima are generally more gradual and widespread as the ocean currents dissipate the hydrothermal plume while the global maximum tends to be very sharp, covering a small region. The gradient ascent strategy finished the fastest of all the strategies.

Future Work

The most important future work is to perform real world tests in well studied areas such as Axial Seamount to further validate the approach. This is the only way to guarantee that the methods are effective. Lacking real world tests, another plume dispersal model, either of a different region or with different plume parameters, could be useful.

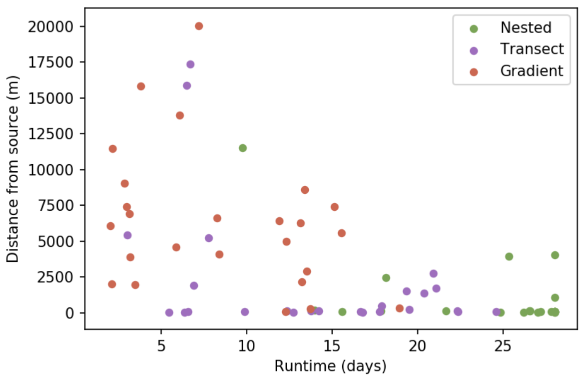

Figure 8: Plot comparing the relationship between between the distance from source and the runtime for each of the runs, colored by strategy. When controlling for the strategy, there is no correlation between the runtime and the distance from source. We can see a cluster for each of the strategies. When not controlling for the strategy, there is a correlation between the runtime and the distance from source.

Currently, the vehicle simulation is rudimentary. Realistic models for sensors such as temperature, optical backscatter, and chemical sensors can be developed. Vehicle resources such as power and data capacity can be implemented. Uncertainty can be added to the vehicles location.

The planning methods have many areas which could use further investigation. The lawnmower surveys could be improved by guaranteeing that potential, but not fully explored, local maxima are verified. Temporal variations in the lateral direction should also be accounted for. This may be particularly important for slower vehicles, perhaps less so if they only move relative to the water, rather than relative to the ground or icy shell. Vehicle resource considerations can be incorporated into the planner. More intelligent path planning can be implemented to reduce resource consumption while performing multiple surveys.

The data volume collected by the vehicle far exceeds the expected communication throughput capabilities. Therefore, a method of summarizing the data collected needs to be developed. A number of spacecraft have implemented systems for this purpose. The Autonomous Sciencecraft Experiment used onboard science algorithms to summarize, delete, and prioritize data for downlink (Chien et al. 2005). The onboard product generation for the Earth Observing-1 mission serves as a predecessor to the proposed HyspIRI Intelligent Payload Module (Chien et al. 2013). The Mars Exploration Rover's (MER) WATCH system processes imagery to detect dust devils and send summarized data products to Earth (Castano et al. 2008). The AEGIS system processes onboard imagery to autonomously retarget science instruments on the Mars Science Laboratory (Francis et al. 2017; Estlin et al. 2014) and MER (Estlin et al. 2012).

Conclusion

We implemented a gradient ascent strategy and a greedy transect strategy for localizing hydrothermal vents. We added a vertical search component to the previously developed nested search strategy. We ran 25 problems per strategy in a hydrothermal simulation environment. The nested and greedy transect strategies frequently located the vent, with the nested strategy finding it most often. The gradient ascent strategy was ineffective at finding the vent. The required search time for the greedy transect strategy was shorter than that of the Nested Bin strategy.

Acknowledgments

The work described by this paper was performed at the Jet Propulsion Laboratory, managed by the California Institute of Technology, under contract to the National Aeronautics and Space Administration.

References

Baker, E. T., and German, C. R. 2004. On the global distribution of hydrothermal vent fields. *Mid-ocean ridges* 245–266.

Baker, E. T.; German, C. R.; and Elderfield, H. 1995. Hydrothermal plumes over spreading-center axes: Global distributions and geological inferences. *Seafloor hydrothermal systems: Physical, chemical, biological, and geological interactions* 47–71.

Branch, A.; Xu, G.; Jakuba, M. V.; German, C. R.; Chien, S.; Kinsey, J. C.; Bowen, A. D.; Hand, K. P.; and Seewald, J. S. 2018. Autonomous nested search for hydrothermal venting. In *Workshop on Planning and Robotics, International Conference on Automated Planning and Scheduling (ICAPS PlanRob 2018)*.

Branch, A.; Flexas, M. M.; Claus, B.; Clark, E. B.; Thompson, A. F.; Chien, S.; Kinsey, J. C.; Fratantoni, D. M.; Zhang, Y.; Kieft, B.; Hobson, B.; and Chavez, F. P. 2019. Front delineation and tracking with multiple underwater vehicles. *Journal of Field Robotics (JFR)* 36(3):568–586.

Branch, A.; McMahon, J.; Xu, G.; Jakuba, M. V.; German, C. R.; Chien, S.; Kinsey, J. C.; Bowen, A. D.; Hand, K. P.; and Seewald, J. S. 2020. Demonstration of autonomous nested search for local maxima using an unmanned underwater vehicle. In *International Conference on Robotics and Automation (ICRA 2020)*.

Burian, E.; Yoerger, D.; Bradley, A.; and Singh, H. 1996. Gradient search with autonomous underwater vehicles using scalar measurements. In *Proceedings of Symposium on Autonomous Underwater Vehicle Technology*, 86–98. IEEE.

Castano, A.; Fukunaga, A.; Biesiadecki, J.; Neakrase, L.; Whelley, P.; Greeley, R.; Lemmon, M.; Castano, R.; and Chien, S. 2008. Automatic detection of dust devils and clouds on mars. *Machine Vision and Applications* 19(5-6):467–482.

Chen, C.; Liu, H.; and Beardsley, R. C. 2003. An unstructured grid, finite-volume, three-dimensional, primitive equations ocean model: Application to coastal ocean and estuaries. *Journal of Atmospheric and Oceanic Technology* 20(1):159–186.

Chien, S., and Wagstaff, K. L. 2017. Robotic space exploration agents. *Science Robotics* 2(7).

Chien, S.; Sherwood, R.; Tran, D.; Cichy, B.; Rabideau, G.; Castano, R.; Davis, A.; Mandl, D.; Trout, B.; Shulman, S.; et al. 2005. Using autonomy flight software to improve science return on earth observing one. *Journal of Aerospace Computing, Information, and Communication* 2(4):196–216.

Chien, S.; Mclaren, D.; Tran, D.; Davies, A. G.; Doubleday, J.; and Mandl, D. 2013. Onboard product generation on earth observing one: a pathfinder for the proposed hyspiri mission intelligent payload module. *IEEE Journal of Selected Topics in Applied Earth Observations and Remote Sensing* 6(2):257–264.

Cruz, N. A., and Matos, A. C. 2010. Adaptive sampling of thermoclines with autonomous underwater vehicles. In *OCEANS 2010*, 1–6. IEEE.

Cruz, N. A., and Matos, A. C. 2014. Autonomous tracking of a horizontal boundary. In *Oceans-St. John's, 2014*, 1–6. IEEE.

Curtin, T. B., and Bellingham, J. G. 2009. Progress toward autonomous ocean sampling networks. *Deep Sea Research Part II: Topical Studies in Oceanography* 56(3):62 – 67. AOSN II: The Science and Technology of an Autonomous Ocean Sampling Network.

Curtin, T. B.; Bellingham, J. G.; Catipovic, J.; and Webb, D. 1993. Autonomous oceanographic sampling networks. *Oceanography* 6(3):86–94.

Das, J.; Py, F.; Maughan, T.; O'Reilly, T.; Messié, M.; Ryan, J.; Sukhatme, G. S.; and Rajan, K. 2012. Coordinated sampling of dynamic oceanographic features with underwater vehicles and drifters. *The International Journal of Robotics Research* 31(5):626–646.

Estlin, T. A.; Bornstein, B. J.; Gaines, D. M.; Anderson, R. C.; Thompson, D. R.; Burl, M.; Castano, R.; and Judd, M. 2012. Aegis automated science targeting for the mer opportunity rover. *ACM Transactions on Intelligent Systems and Technology (TIST)* 3(3):50.

Estlin, T.; Gaines, D.; Bornstein, B.; Schaffer, S.; Tompkins, V.; Thompson, D. R.; Altinok, A.; Anderson, R. C.; Burl, M.; Castaño, R.; et al. 2014. Automated targeting for the msl rover chemcam spectrometer. In *12th International Symposium on Artificial Intelligence, Robotics, and Automation in Space (i-SAIRAS)*, 17–19.

Farrell, J. A.; Pang, S.; and Li, W. 2005. Chemical plume tracing via an autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering* 30(2):428–442.

Ferri, G.; Jakuba, M. V.; and Yoerger, D. R. 2010. A novel trigger-based method for hydrothermal vents prospecting using an autonomous underwater robot. *Autonomous Robots* 29(1):67–83.

Flexas, M. M.; Troesch, M. I.; Chien, S.; Thompson, A. F.; Chu, S.; Branch, A.; Farrara, J. D.; and Chao, Y. 2018. Autonomous sampling of ocean submesoscale fronts with ocean gliders and numerical model forecasting. *Journal of Atmospheric and Oceanic Technology* 35(3):503–521.

Francis, R.; Estlin, T.; Doran, G.; Johnstone, S.; Gaines, D.; Verma, V.; Burl, M.; Frydenvang, J.; Montaño, S.; Wiens, R. C.; Schaffer, S.; Gasnault, O.; DeFlores, L.; Blaney, D.; and Bornstein, B. 2017. Aegis autonomous targeting for chemcam on mars science laboratory: Deployment and results of initial science team use. *Science Robotics* 2(7).

German, C., and Seyfried, W. 2014. Hydrothermal processes. *Treatise on geochemistry* 8:191–233.

German, C. R.; Yoerger, D. R.; Jakuba, M.; Shank, T. M.; Langmuir, C. H.; and Nakamura, K.-i. 2008. Hydrothermal exploration with the autonomous benthic explorer. *Deep Sea Research Part I: Oceanographic Research Papers* 55(2):203–219.

Godin, M. A.; Zhang, Y.; Ryan, J. P.; Hoover, T. T.; and Bellingham, J. G. 2011. Phytoplankton bloom patch center localization by the tethys autonomous underwater vehicle. In *OCEANS'11 MTS/IEEE KONA*, 1–6.

Haley, P.; Lermusiaux, P.; Robinson, A.; Leslie, W.; Logoutov, O.; Cossarini, G.; Liang, X.; Moreno, P.; Ramp, S.; Doyle, J.; Bellingham, J.; Chavez, F.; and Johnston, S. 2009. Forecasting and reanalysis in the monterey bay/california current region for the autonomous ocean sampling network-ii experiment. *Deep Sea Research Part II: Topical Studies in Oceanography* 56(3):127 – 148. AOSN II: The Science and Technology of an Autonomous Ocean Sampling Network.

Hand, K. P.; Chyba, C. F.; Priscu, J. C.; Carlson, R. W.; and Nealson, K. H. 2009. Astrobiology and the potential for life on europa. In Pappalardo, R. T.; McKinnon, W. B.; and Khurana, K., eds., *Europa*. University of Arizona Press. 589–630.

Hsu, H.-W.; Postberg, F.; Sekine, Y.; Shibuya, T.; Kempf, S.; Horányi, M.; Juhász, A.; Altobelli, N.; Suzuki, K.; Masaki, Y.; et al. 2015. Ongoing hydrothermal activities within enceladus. *Nature* 519(7542):207.

Jakuba, M., and Yoerger, D. R. 2008. Autonomous search for hydrothermal vent fields with occupancy grid maps. In *Proc. of ACRA*, volume 8, 2008.

Jia, X.; Kivelson, M. G.; Khurana, K. K.; and Kurth, W. S. 2018. Evidence of a plume on europa from galileo magnetic and plasma wave signatures. *Nature Astronomy* 2:459–464.

Jordan, S. F.; Rammu, H.; Zheludev, I. N.; Hartley, A. M.; Marechal, A.; and Lane, N. 2019. Promotion of protocell self-assembly from mixed amphiphiles at the origin of life. *Nature ecology & evolution* 3(12):1705–1714.

Leonard, N. E.; Paley, D. A.; Lekien, F.; Sepulchre, R.; Fratantoni, D. M.; and Davis, R. E. 2007. Collective motion, sensor networks, and ocean sampling. *Proceedings of the IEEE* 95(1):48–74.

Nakamura, K.; Veirs, S.; Sarason, C. P.; McDuff, R. E.; Stahr, F.; Yoerger, D. R.; and Bradley, A. M. 2000. Electrochemical signals in rising buoyant plumes and tidally oscillating plumes at the main endeavour vent field, juan de fuca ridge. *EOS, Transactions of the American Geophysical Union* 81(48).

National Aeronautics and Space Administration. 2018. Ocean worlds.

Pang, S. 2010. Plume source localization for auv based autonomous hydrothermal vent discovery. In *OCEANS 2010*, 1–8. IEEE.

Petillo, S.; Schmidt, H.; and Balasuriya, A. 2012. Constructing a distributed auv network for underwater plume-tracking operations. *International Journal of Distributed Sensor Networks* 2012:Article ID 191235, 12pp.

Ramp, S.; Davis, R.; Leonard, N.; Shulman, I.; Chao, Y.; Robinson, A.; Marsden, J.; Lermusiaux, P.; Fratantoni, D.; Paduan, J.; Chavez, F.; Bahr, F.; Liang, S.; Leslie, W.; and Li, Z. 2009. Preparing to predict: The second autonomous ocean sampling network (aosn-ii) experiment in the monterey bay. *Deep Sea Research Part II: Topical Studies in Oceanography* 56(3):68 – 86. AOSN II: The Science and Technology of an Autonomous Ocean Sampling Network.

Saigol, Z.; Dearden, R.; Wyatt, J.; and Murton, B. 2010. Belief change maximisation for hydrothermal vent hunting using occupancy grids. In *Proceedings of the Eleventh Conference Towards Autonomous Robotic Systems (TAROS-10)*, 247–254.

Sparks, W. B.; Hand, K.; McGrath, M.; Bergeron, E.; Cracraft, M.; and Deustua, S. 2016. Probing for evidence of plumes on europa with hst/stis. *The Astrophysical Journal* 829(2):121.

Sparks, W. B.; Schmidt, B. E.; McGrath, M. A.; Hand, K. P.; Spencer, J. R.; Cracraft, M.; and Deustua, S. E. 2017. Active cryovolcanism on europa? *The Astrophysical Journal* 839(2):L18.

Speer, K. G., and Rona, P. A. 1989. A model of an atlantic and pacific hydrothermal plume. *Journal of Geophysical Research: Oceans* 94(C5):6213–6220.

Sun, L.; Li, Y.; Yan, S.; Wang, J.; and Chen, Z. 2016. Thermocline tracking using a portable autonomous underwater vehicle based on adaptive threshold. In *OCEANS 2016-Shanghai*, 1–4. IEEE.

Tian, Y.; Zhang, A.; Li, W.; Yu, J.; Li, Y.; and Zeng, J. 2014. A behavior-based planning strategy for deep-sea hydrothermal plume tracing with autonomous underwater vehicles. In *OCEANS 2014-TAIPEI*, 1–10. IEEE.

Trumbo, S. K.; Brown, M. E.; and Hand, K. P. 2019. Sodium chloride on the surface of europa. *Science Advances* 5(6).

Turner, J. S. 1979. *Buoyancy effects in fluids*. Cambridge University Press.

Veirs, S. R. 2003. *Heat flux and hydrography at a submarine volcano: Observations and models of the Main Endeavour vent field in the northeast Pacific*. Ph.D. Dissertation, University of Washington.

Waite, J. H.; Glein, C. R.; Perryman, R. S.; Teolis, B. D.; Magee, B. A.; Miller, G.; Grimes, J.; Perry, M. E.; Miller, K. E.; Bouquet, A.; Lunine, J. I.; Brockwell, T.; and Bolton, S. J. 2017. Cassini finds molecular hydrogen in the enceladus plume: Evidence for hydrothermal processes. *Science* 356(6334):155–159.

Yoerger, D. R.; Bradley, A. M.; Jakuba, M. V.; Tivey, M. A.; German, C. R.; Shank, T. M.; and Embley, R. W. 2007a. Mid-ocean ridge exploration with an autonomous underwater vehicle.

Yoerger, D. R.; Jakuba, M.; Bradley, A. M.; and Bingham, B. 2007b. Techniques for deep sea near bottom survey using an autonomous underwater vehicle. *The International Journal of Robotics Research* 26(1):41–54.

Zhang, Y.; Bellingham, J. G.; Godin, M.; Ryan, J. P.; McEwen, R. S.; Kieft, B.; Hobson, B.; and Hoover, T. 2010. Thermocline tracking based on peak-gradient detection by an autonomous underwater vehicle. In *OCEANS 2010*, 1–4. IEEE.

Zhang, Y.; Bellingham, J. G.; Ryan, J. P.; Kieft, B.; and Stanway, M. J. 2013. Two-dimensional mapping and tracking of a coastal upwelling front by an autonomous underwater vehicle. *Proc. MTS/IEEE Oceans'13* 1–4.

Zhang, Y.; Bellingham, J. G.; Ryan, J. P.; Kieft, B.; and Stanway, M. J. 2016. Autonomous four-dimensional mapping and tracking of a coastal upwelling front by an autonomous underwater vehicle. *Journal of Field Robotics* 33(1):67–81.

Zheng, L., and Weisberg, R. H. 2012. Modeling the west florida coastal ocean by downscaling from the deep ocean, across the continental shelf and into the estuaries. *Ocean Modelling* 48:10 – 29.