

Autonomous Search for Hydrothermal Vent Fields with Occupancy Grid Maps

Michael V. Jakuba

Australian Centre for Field Robotics
University of Sydney
m.jakuba@acfr.usyd.edu.au

Dana R. Yoerger

Applied Ocean Physics and Engineering
Woods Hole Oceanographic Institution
dyoerger@whoi.edu

Abstract

This paper demonstrates the application of occupancy grid mapping to autonomous robotic chemical plume source localization. We formulate the chemical plume tracing problem as one of resolving an occupancy grid map of candidate plume source locations. Our method’s significance is its applicability to scenarios with multiple plume sources. We demonstrate the method on hydrothermal tracer data collected by an Autonomous Underwater Vehicle (AUV) in the field. The resultant maps indicate both coverage and likely vent locations.

1 Introduction

Potential robotic applications for chemical plume tracing include [Russell, 2001] pollution and environmental monitoring, chemical plant safety, search and rescue (facilitated by tracking carbon dioxide plumes emitted by survivors), explosive ordinance removal, and deep sea hydrothermal vent prospecting. Chemical plume tracing algorithms that have been implemented on robots or designed for them range from concentration gradient ascent, to biologically inspired algorithms [Farrell *et al.*, 2005; Grasso *et al.*, 2000; Kuwana *et al.*, 1995; Nagasawa *et al.*, 1999; Ishida *et al.*, 2001; Consi *et al.*, 1994; Lilienthal and Duckett, 2003a; Balkovsky and Shraiman, 2002], to model-based strategies that estimate source location along with other plume parameters [Ishida *et al.*, 1998; Christopoulos and Roumeliotis, 2005a] or that build maps of probable source location [Pang and Farrell, 2006; Lilienthal and Duckett, 2003b], with some work also exploring multi-agent cooperative approaches either biologically-inspired [Hayes *et al.*, 2002; Ferri *et al.*, 2006] or model-based but with the advantage of distributed sensing [Christopoulos and Roumeliotis, 2005b; Zarzhitsky *et al.*, 2004]. Kowadlo and Russell [2008] provide an extensive overview.

Existing solutions do not adequately address a key aspect common to many real-world applications: an

initially unknown number of potential plume sources. When an unknown number of sources (possibly zero) exist within a prescribed area, the task becomes one of finding them all and, most importantly, confidently declaring that no further unlocalized sources exist in the environment. For instance, a robot monitoring a chemical plant for leaks should not only be able to rapidly localize a leak once detected, but should also actively avoid resampling the same potential leak sites in order to conduct its patrol as efficiently as possible. For routine leak monitoring, the most likely outcome is that no plume sources exist—a conclusion that should be reported with an associated notion of its certainty. Other scenarios intrinsically involve a search for multiple plume sources: A navy may wish to be confident a harbor is completely free of mines. Likewise, the search for survivors following the collapse of a building should proceed until no reasonable chance exists to find additional plumes of respired carbon dioxide.

These examples illustrate that when an initially unknown number of chemical sources might be present in the survey domain, it is often important to find them all. The aspect common to both inherently multi-source scenarios and scenarios like routine leak detection monitoring, where zero or only a single source will typically exist, is a requirement to assess coverage. Probabilistic assessments of coverage are key to ascertaining what portion of an environment has in fact been adequately searched. Coverage, if computable in real time, also makes adaptive exploration possible and could yield many-fold improvements in survey time over dense regular sampling.

Hydrothermal vent prospecting, the chemical plume tracing application studied here, consists of inferring the location of sea floor vent sites from water-column measurements of the emitted plumes. It is an inherently exploratory task, and assessing coverage is therefore a key component. Furthermore, multiple vent fields often occur within hundreds of meters of one another on geologically active sections of sea floor rift valleys. Our map-based approach to chemical plume source localiza-

tion was developed to satisfy both requirements.

2 Approach

Our approach to chemical plume source localization utilizes occupancy grid (OG) mapping [Elfes, 1989; Moravec and Elfes, 1985] to create maps of likely source locations and, critically, also to indicate locations unlikely to contain sources. The latter aspect enables maps produced using our method to indicate coverage, a property possessed by cellular decompositions of the environment, of which OGs are but one example [Choset, 2001]. We redefine the binary state of grid cells to denote the presence or absence of an active chemical source. Because OG maps explicitly represent empty space with a degree of confidence, maps can be used to assess not only probable source locations but also whether a survey area has been adequately searched, i.e. whether there are likely to be any additional sources.

OG maps were originally developed to integrate successive measurements from low-cost but noisy sonar range finders¹ into a consistent map encapsulating the robot’s uncertainty about its environment [Martin and Moravec, 1996]. In our approach, the plumes themselves become the sensors used to infer the locations of their sources. Plume detections imply nearby or upwind sources, but non-detections also serve to constrain source location by reducing the likelihood that a source lies nearby or upwind. Since the spatial relationship between where a plume is detected and the location of its source is inherently uncertain, using plumes as noisy sensors demands a stochastic mapping approach.

Both Pang et al. [Pang and Farrell, 2006] and Farrell et al. [Farrell et al., 2003] use model-based probabilistic reasoning to learn source location on a grid-based representation of the environment akin to that employed here. Our method owes its inspiration to Bayesian methodology first described in Pang’s PhD work [Pang, 2004] for learning the state of a binary random field of the possible locations for a plume source. Our work represents a significant extension that applies to environments with an unknown number of sources. The difference is significant because assuming a single source represents a powerful constraint on the state of the environment that enables inferences to be made about the environment in its entirety regardless of the portion of the environment actually observed. Consequently, maps built assuming a single source do not indicate coverage.

Our approach necessarily relies on a model for plume behavior. Model-based approaches to robotic chemical plume tracing that rely on measured concentration are

¹In fact, much of the “noise” in data from these sensors can be attributed to deterministic physical phenomena [Leonard and Durrant-Whyte, 1992].

prone to performance limitations that stem from the slow convergence, in turbulent plumes, of statistics based on mean concentration [Liao and Cowen, 2002]. Our approach does not attempt to infer anything from measured concentration, and instead relies exclusively on binary detections and non-detections of plume effluent. Following the work of Farrell and Pang et. al. [Farrell et al., 2003; Pang and Farrell, 2006] we model the likelihood of plume detection as a function of advection and the relative location between the source and detector. In isolation, such measurements are sufficient only to constrain source location to lie at an unknown distance upwind.² The OG map provides the mechanism for fusing multiple measurements from different locations and times into consistent estimates of source locations. We note that while models are necessary to create maps, mapping itself does not preclude the use of biomimetic or other trajectories, and furthermore [Pang and Farrell, 2006] that maps are useful especially when other search algorithms fail.

A defining characteristic of chemical plume source localization considered within the context of occupancy grid mapping is the small number of cells expected to be occupied by active plume sources. From a Bayesian perspective, the prior probability of occupancy associated with each cell is much lower than is typical in OG maps constructed from laser or sonar range-finder data (e.g. [Thrun et al., 2005]). In this work we employ a map update rule [Jakuba, 2007] different from the classical Bayesian map update rule described in [Elfes, 1989]. The maps produced by our algorithm, in the context of multiple plume source localization, tend to identify cells likely to contain sources more efficiently than the classical rule and produce maps that are more consistent with the true number of occupied cells.³ The computational cost is equivalent to the classical algorithm and suitable for real time implementation.

3 Algorithm

3.1 The classical Bayesian Algorithm

Classical Bayesian OG mapping seeks to compute the marginal posterior probabilities of occupancy $p(\mu_c = 1 | Z^t)$ for each grid cell, where $\mu_c \in \{0, 1\}$ denotes the emptiness $\mu_c = 0$ or occupancy $\mu_c = 1$ of cell number c , and $Z^t = \{z^1, z^2, \dots, z^t\}$ denotes the set of sensor measurements up to time t . For computational and storage reasons,⁴ the marginal posteriors are com-

²Wind history, rather than just instantaneous wind must be taken into account in variable winds.

³Quantitative inter-algorithm comparisons appear in [Jakuba, 2007; Ferri et al., 2007].

⁴The number of possible maps for a grid with C cells is 2^C .

puted instead of the full posterior over all possible maps, $p(\mu_1, \mu_2, \dots, \mu_C \mid Z^t)$.

Unfortunately, it is not in general possible to arrive at exact marginal posteriors without first computing the full posterior and then carrying out the marginalization. To avoid both the computational and storage burden implied by this task, typical OG map update rules require assumptions that enable direct computation of the marginal posteriors from the measurements. Most update rules are also recursive, for the sake of real-time operation.

The classical Bayesian update rule produces inexact marginal posteriors but it is simple to implement and popular (e.g. the numerous examples of successful applications in Ch. 9 of [Thrun *et al.*, 2005]). To simplify notation let m_c and \bar{m}_c denote the events that cell c is occupied $\mu_c = 1$ or empty $\mu_c = 0$, respectively. The update rule is applied following each sensor measurement separately to each cell within the perceptual field of the sensor:

$$r_c^t = \frac{P[m_c \mid z^t]}{1 - P[m_c \mid z^t]} \cdot \frac{1 - P[m_c]}{P[m_c]} \cdot r_c^{t-1} \quad (1)$$

where $r_c^t \triangleq P[m_c \mid Z^t] / P[\bar{m}_c \mid Z^t]$ denotes the posterior odds of occupancy (the ‘‘odds ratio’’), and from which the posterior probability of occupancy can be trivially recovered. All that is required to implement this algorithm is a prior probability of occupancy for each cell $P[m_c]$ and a means of computing a single-cell inverse sensor model $P[m_c \mid z^t]$ based on sensor characteristics and the relative position between each cell and the robot.

3.2 A Sensor Model for Plume Source Localization

A characteristic common to all plume source localization scenarios is that measurements fundamentally consist of plume detections and non-detections. Ancillary information like chemical concentration, wind direction, and other measurements can constrain the likely origins of detected plumes, or in the case of non-detections suggest what regions of a map are likely to be empty. These measurements are therefore important, but they are not fundamental in the sense that the set of all measurements can be classified into two non-intersecting sets, one of detections and the other of non-detections. The state space consists of plume source locations (here the binary state of each grid cell) which does not include the state of the actual plumes emitted from them. Consequently, the plumes themselves (and the measurements that observe their state) can be regarded as part of a composite dynamic sensor analogous to whiskers that indicate the presence or absence of an obstacle at locations that depend on whisker orientation. In parallel with the notation for the binary state of a grid cell, let d^t and \bar{d}^t

denote the events of a detection $z^t = 1$ or non-detection $z^t = 0$, respectively.

A simple model for the probability of not detecting a plume given knowledge of the locations of all sources (i.e. the true state of each grid cell in the map) is:

$$P[\bar{d}^t \mid \mu_1, \mu_2, \dots, \mu_C] = (1 - P_F) \prod_{c:m_c}^C (1 - P_c^t), \quad (2)$$

where P_F denotes the probability of a false detection and P_c^t denotes the probability that sufficient signal from a source in cell c arrived at the detector at time t to trigger a detection. The probability of registering a detection $P[d^t \mid \mu_1, \mu_2, \dots, \mu_C]$ is one minus this quantity.

The sensor model described by (2) is a forward model. That is, it describes the likelihood of a measurement given a known map. The parameters of forward models, i.e. the P_c^t and P_F in (2), are usually easy to acquire from calibration data [Thrun, 2003] or in this case to predict based on a plume model.

The classical OG mapping update rule requires the marginalized inverse of the quantity given by (2). In fact, (2) is both invertible and marginalizable. The result [Jakuba, 2007] is

$$P[m_c \mid d^t] = \frac{1 - (1 - P_F)(1 - P_c^t) \prod_{s \neq c}^C (1 - P_s^t P[m_s])}{1 - (1 - P_F) \prod_{s=1}^C (1 - P_s^t P[m_s])} P[m_c] \quad (3a)$$

$$P[m_c \mid \bar{d}^t] = \frac{1 - P_c^t}{1 - P_c^t P[m_c]} P[m_c]. \quad (3b)$$

Note that the large sums over possible maps $\{\mu : m_c\}$ implied by direct application of Bayes Rule to (2) have been replaced with products over the number of cells in the map C . These equations also provide the key to the alternative map update rule employed in this work.

3.3 An Alternative Map Update Rule

For plume source localization problems where sensor models of the form (2) apply, and where the expected number of sources (occupied cells) is small but unknown, an alternative map update rule often produces considerably better results than the classical update rule.

The marginal posteriors given by the inverse sensor model (3) are exact. Therefore, if the set of all measurements consists of a single measurement, i.e. if $Z^t = \{z^1\}$, then (3) can be used to compute the exact marginal posteriors. Now suppose these one-step marginal posteriors are regarded as independent, i.e., $p(\mu_1, \mu_2, \dots, \mu_C \mid Z^1) = \prod_{c=1}^C (P[\mu_c \mid Z^1])$. This same requirement is typically placed on the prior $p(\mu_1, \mu_2, \dots, \mu_C)$, thus posteriors satisfying this property can be regarded as specifying a new, ‘‘revised,’’ prior. Furthermore, if at $t = 2$ the priors $P[m_c]$ in (3) are replaced by the single-step marginal posteriors

$P[m_c | Z^1]$, and if these are independent, then the resultant marginal posteriors $P[m_c | Z^2]$ produced will remain exact.

Replacing the priors in (3) with the revised priors $\tilde{P}_c^t \triangleq P[m_c | Z^{t-1}]$ at each time step defines an alternative map update rule. The new rule expressed in terms of the odds ratio r_c^t is:

$$r_c^t = \begin{cases} \frac{1 - (1 - P_F^t)(1 - P_c^t) \prod_{s \neq c} (1 - P_s^t \tilde{P}_s^t)}{1 - (1 - P_F^t) \prod_{s \neq c} (1 - P_s^t \tilde{P}_s^t)} \cdot r_c^{t-1}, & \text{if } d^t \quad (4a) \\ (1 - P_c^t) r_c^{t-1}, & \text{if } \bar{d}^t \quad (4b) \end{cases}$$

where the revised prior can be computed from the odds ratio as $\tilde{P}_c^t = \frac{r_c^{t-1}}{1 + r_c^{t-1}}$. Unlike (1), this alternative update rule cannot be written as a single equality because its simple multiplicative structure depends on the particular form of the inverse sensor model (3). Because OG algorithms intrinsically involve conditional independence assumptions, we have named our algorithm after its particular assumption: Independence of Posteriors (IP).

In principle, independence of the posteriors could be assumed in conjunction with any sensor model; however, the forward model (2) is special because the priors $\tilde{P}_c^0 \triangleq P[m_c]$ enter explicitly into the inverse sensor model generated from it, thereby providing a mechanism for interpreting the current measurement in terms of the present belief in the states of the cells in the map. Like other “context-sensitive” OG methods specific to sonar range finders [Moravec and Cho, 1989; Lim and Cho, 1992], the IP algorithm interprets new measurements in light of the present state of the map.

4 Field Results: Hydrothermal Vent Localization

OG mapping is applied in this section to field data from deep sea hydrothermal vent prospecting missions carried out by an autonomous underwater vehicle (AUV). OG mapping represents a powerful data processing technique for condensing multi-modal sensory data into a simple representation of likely sea floor vent locations. The maps produced indicate both coverage and likely vent field locations. As such they are valuable tools for humans planning follow-on surveys, but the simplicity of the representation holds out the promise of fully autonomous interpretation and retasking.

4.1 Hydrothermal Vent Prospecting

Deep sea hydrothermal vent fields occur throughout the 60,000 km length of the worldwide subsea mountain chain known as the Mid-Ocean Ridge. They occur where volcanic and tectonic processes induce the circulation of sea water through young oceanic crust. Hot, chemically altered sea water is persistently discharged, or vented,

as a turbulent, buoyant plume that rises above the sea floor to heights of 100 m to 400 m in typical hydrographic settings [Speer and Rona, 1989], and then spreads laterally along isopycnals in essentially a smog layer known as the non-buoyant plume. This structure suggests a natural strategy for homing on the source of venting: Establish contact with the large spatial signature of the non-buoyant plume; find the buoyant stems within, and finally follow these to the sea floor.

Recent expeditions utilizing the Autonomous Benthic Explorer (ABE) AUV [Yoerger *et al.*, 1991] have demonstrated the effectiveness of an AUV for vent localization [German *et al.*, 2008]. On these expeditions, the AUV (Fig. 1) performed three successive nested survey dives at each site, according to the methodology detailed in [German *et al.*, 2008]. Each stage was executed at progressively lower altitudes and finer trackline spacing based on the results of the previous stage and any available ancillary data, a process that broadly reflects the physical structure of hydrothermal plumes themselves.

The AUV was recovered between each stage in order to analyze sensor data and plan the next stage. It is a testament to the expertise of the scientific staff on these expeditions that this process was largely successful; however, it also was often contentious and potentially much less efficient than in situ autonomous retasking. Autonomous retasking would eliminate the need for vehicle recovery and redeployment between stages. OG mapping provides a mechanism to condense multi-modal sensor data via encapsulated expert knowledge in the form of models for plume behavior into a map of the sea floor colored by information ideal for planning follow-on surveys. Maps show likely vent field locations along with coverage—an indication of whether other, undetected vent fields are likely to be found in the survey area. This is important information as it is common for several vent fields to exist on the same ridge segment within hundreds of meters of one another. The maps are already valuable to human mission planners, but because they can be computed and interpreted autonomously in real time, they promise to enable AUVs to carry out hydrothermal vent prospecting autonomously.

4.2 Implementation

The OG maps presented in the next section were generated from automatic classification of plume tracer data into binary detections and non-detections of buoyant plume effluent. Buoyant hydrothermal plumes occupy smaller spatial and temporal scales than occupied by non-buoyant hydrothermal plumes and buoyant plume detections are consequently easier to interpret in terms of likely source location.



Figure 1: The ABE AUV being deployed on a hydrothermal vent prospecting mission in 2005 on the Southern Mid-Atlantic Ridge near Ascension Island. Various acoustic sensors in the lower pod measure vehicle velocity over ground and through water, collect soundings of sea floor depth, and collect rangings to fixed acoustic beacons. Hydrographic sensors located near the nose of the starboard and lower pods collect physical oceanographic data and record the concentration of various hydrothermal tracers.

Buoyant plume dynamics The dynamics of typical buoyant hydrothermal plumes are dominated by turbulent entrainment [Turner, 1986] of the surrounding sea water. This process dilutes vented hydrothermal effluent by a factor of 10^4 – 10^5 by the time the density of the mixed fluid matches that of the surrounding fluid at its terminal rise height [Lupton *et al.*, 1985]. Turbulent entrainment of ambient waters makes regional hydrography an important influence on the trajectory followed by an ascending buoyant plume (BP), and acts directly to obscure source location by expanding plume diameter during ascent. Ambient crossflows also influence plume trajectory by bending the plume centerline in the direction of flow [Middleton and Thomson, 1986].

Forward sensor model In our approach, the physics of plume dispersal are encapsulated within the factor P_c^t from (2) that describes the likelihood of encountering effluent from a hypothetical source in cell c at a time-

dependent robot location. To generate an analytical form for P_c^t , we employed field data from over 10 deep-sea deployments where both groundtruth vent locations and water-column hydrothermal tracer data were available. Jakuba [2007] contains the details; what follows is an overview.

Buoyant hydrothermal plumes are three-dimensional structures, with motion in the vertical dimension driven by the buoyancy of the plume itself. Self-generated turbulence (the energy being supplied by the plume’s own buoyancy) results in relatively well mixed plumes compared to those typically studied in robotic chemical plume tracing applications. Nevertheless, synoptic concentration profiles still deviate strongly from empirically observed time-averaged profiles [Lupton, 1995]. Individual vents tend to be clustered in fields so that a vent field as a whole produces a collection of plumes, some of which coalesce into larger plumes.

Our model relates the probability of encountering hydrothermal effluent near a vent field to vehicle position \mathbf{x}_v , vent position \mathbf{x}_s , and vent-vehicle vertical separation h_{vs} :

$$P_c^t = 1 - \left(1 - \left(b_0 + \frac{1}{2}ah_{vs} \right)^2 \cdot \frac{1}{2\pi \left(\left(\frac{1}{2}ah_{vs} \right)^2 + \sigma_s^2 \right)} \right. \\ \left. \cdot \exp \left(- \frac{|\mathbf{x}_s + \mathbf{U}_0 \frac{h_{vs}}{W_0} - \mathbf{x}_v|^2}{2 \left(\left(\frac{1}{2}ah_{vs} \right)^2 + \sigma_s^2 \right)} \right) \right)^Q \quad (5)$$

where b_0 and σ_s encode characteristic source sizes for individual vents and entire vent fields respectively, \mathbf{U}_0 denotes the average horizontal water velocity over the course of plume rise to survey height, a encapsulates the effect of plume growth via turbulent entrainment, and Q denotes the number of discrete plumes emanating from a field.⁵ The effect of buoyancy is accommodated by a deterministic rise rate of $W_0 = 0.1$ m/s, the depth-averaged value suggested in [Speer and Rona, 1989]. The localization performed applies to entire vent fields rather than to individual vents.

The model encapsulates two important effects: (1) advection by background crossflows, and (2) rise-height dependence. Plumes emanating from a vent field are more likely to be detected at a given altitude near their time-averaged centerline, i.e. when $|\mathbf{x}_s + \mathbf{U}_0 \frac{h_{vs}}{W_0} - \mathbf{x}_v|$ is small. Increased vertical vent-vehicle separation creates wider plumes that are more likely to be intercepted by a passing vehicle but that also contain less information about their sources’ locations. That tradeoff is encapsulated by the height-dependent factors, $b_0 + \frac{1}{2}ah_{vs}$ (plume width), and $\left(\frac{1}{2}ah_{vs} \right)^2 + \sigma_s^2$ (random-walk scale).

⁵This information is obviously unavailable a priori. We have attained satisfactory results with $Q = 5$ in a variety of geological settings in both the Atlantic and Pacific oceans.

Buoyant plume detection Hydrothermal effluent contains a wide variety of chemical species that react at varying rates with sea water [Lilley *et al.*, 1995]. In our approach, detecting buoyant plume effluent depended on measurements of non-conservative tracers, tracers whose concentrations decayed rapidly away from plume sources due to chemical reactions or biological processes in addition to passive advection and diffusion. The relatively small spacial signature of buoyant plumes in comparison to the survey area makes outlier detection methods a good fit for detection. In our experiments we decorrelated each data stream by downsampling, then batch-processed each data stream with a Hampel Identifier [Davies and Gather, 1993] followed by a consensus vote between tracers suited to detection at survey height [Jakuba, 2007].

In addition to measuring hydrographic tracers, the ABE AUV makes two other measurements in support of hydrothermal vent prospecting operations: (1) its position within an acoustic transponder net, and (2) water current velocity. Position measurements allow plume data to be georegistered, an essential requirement for building a consistent OG map. Water current velocity measurements allow the locations of samples to be back-propagated to their origin on the sea floor.

Grid parameters The sea floor defines the surface on which vent fields must lie. Because rise height affects the information content of buoyant plume detections, accurate knowledge of vehicle-sea floor separation is required. In practice, bathymetry acquired acoustically from a surface ship with order 100 m horizontal resolution is adequate. Our forward model applies to the locations of entire vent fields. Grid cell size should be small relative to the characteristic uncertainty associated with source location (about 25 m in the data sets studied) to avoid discretization artifacts. For high altitude surveys, a lower resolution is sufficient and also desirable because these dives typically span a larger area and computational load grows quadratically with survey area (linearly with the number of grid cells).

An appropriate cell prior reflects the expected number of vent fields in the survey area scaled by grid resolution. Let ρ_f denote the expected the vent field density (vents/unit area). For the expected number of occupied cells to match the prior probability of occupancy, $P[m_c]$, is set to

$$P[m_c] = \Delta^2 \rho_f, \quad (6)$$

where Δ denotes grid cell dimension. The OG maps presented in the next section were initialized with an assumed vent field density of $\rho_f = 1 / \text{km}^2$, yielding priors between 10^{-5} and 10^{-3} depending on grid cell size.




Figure 2: OG Map produced from dive ABE-128 (250 m altitude). The vehicle registered detections near two of the three vent fields ultimately localized in addition to a detection near the southern end of the survey, the source location of which remains unknown.

4.3 Results

Figures 2 and 3 show OG maps generated by the IP OG algorithm from data collected on two separate dives, ABE-128 (250 m altitude) and ABE-136 (the subsequent 50 m altitude dive). Regions with relatively high posterior odds of occupancy (hot/light shading) agree well with clusters of groundtruth vent locations. Regions with relatively low posterior odds (cool/dark shading) indicate coverage over regions where no buoyant plumes were encountered.⁶

Lowered posterior odds in Fig. 2 trace out an image of the AUV’s track on sea floor with a relative spacing dependent on the direction and magnitude of water current velocity at the sampling time. The trace indicates track-line spacing was insufficiently dense by at least a factor of two. A subsequent survey designed purely based upon this map would have discovered two of the three vent

⁶Because the perceptual fields associated with detections from disparate vent fields in these maps do not intersect and are few in number, it is possible to compute the exact marginal posteriors. The result [Jakuba, 2007] is similar to the results shown here but is much more expensive to compute.

fields shown and possibly a third site to the south that was never groundtruthed. The survey would have missed the field centered at (7.8, 7.8); however, that is consistent with the field’s location outside the region indicated as covered in the OG map.

For nested surveys, the finished OG map from one stage can be used to initialize a higher resolution OG map for the subsequent stage. Figure 3 illustrates the result of using the posterior from a previous dive as a prior. ABE-128 and ABE-136 were conducted at ~ 250 m and 50 m above bottom respectively. Because of this difference in height, the grid cell sizes chosen were different: 20 m and 5 m on a side respectively. Each cell in the map shown in Fig. 2 was split into sixteen descendant cells, each of which was initialized with a probability of occupancy $1/16$ that of the posterior in the original parent cell.⁷ To do so requires regarding the final posteriors of the original map as independent, an assumption which is false but consistent with the assumption required by the IP algorithm.

Detections at the southern edge of the survey shown in Fig. 3 corroborate evidence from the high-altitude (250 m) dive at this site (Fig. 2 for a southern vent field; however, this southern site was never groundtruthed). Comparing low altitude (5 m) coverage (shown as white boxes in Fig. 3) with the OG map provides a likely explanation. The map indicates the southern 5 m altitude survey was misplaced, and should have been centered some 200 m to the south. The 5 m altitude survey was designed while at sea on the original expedition, and before we had applied OG mapping to hydrothermal vent prospecting. Likewise the central portion of the OG map indicates that better coverage over likely vent sites could have been attained by decreasing the north-south extent of the northern survey in favor of extending it to the west. The three western groundtruth vent locations that lie outside the bounds of the 5 m altitude survey were discovered by a remotely operated vehicle (ROV) several months after the completion of ABE operations at this site but lie within the areas indicated in the OG map as likely to contain vents.

5 Conclusion

In this work, we redefined the binary state of grid cells in an OG map to denote either the presence (occupancy) or absence (vacancy) of an active plume source. The result is an algorithm capable of localizing plume sources in environments where the number of sources is unknown a priori. OG mapping applied to plume source localization enables both plume detections and plume absence to inform the estimated state of the map. The result

⁷Kraetzschmar et al. [Kraetzschmar *et al.*, 2004] use this same “probabilistic mean” to propagate probabilities of occupancy through a multi-resolution probabilistic quad tree.

indicates probable source locations and explicitly represents explored but empty portions of the domain. The latter property indicates coverage, which is important in multiple-source scenarios, for example to direct search or declare search within an area complete.

Application to hydrothermal plume data collected by the ABE AUV during hydrothermal vent prospecting operations in the deep ocean demonstrated the applicability of our methods to real-world scientific exploration. The automatically generated OG maps condense multimodal sensory data into easily interpreted maps useful for planning nested surveys. The OG maps can also be computed in real time, potentially facilitating the autonomous design of nested surveys or more general survey trajectories and obviating the need for human data




Figure 3: OG Map produced from dive ABE-136 (50 m altitude) using the map of Fig. 2 to define the prior. The locations of all vent fields ultimately confirmed agree well with the regions of the map populated by cells with high posterior odds (hot/light shading). Regions with low posterior odds (cool/dark shading) indicate fairly uniform coverage over most of the survey extent, though with some patchiness particularly in the southern half of the survey area.

review and low-level survey design.

Acknowledgment

We are grateful to Drs. C. German, T. Shank, and C. Langmuir for providing opportunities to extend ABE's capabilities at sea and to Dr. K. Nakamura for graciously loaning us an Eh probe of his design since 1999, without which detecting buoyant hydrothermal plumes would have been considerably more difficult. Our thanks go to the ABE team and to the captains and crews of the R/V Kilo Moana and the RRS Charles Darwin for their critical role in generating the ABE operations results herein. The authors gratefully acknowledge the support of NSF (OCE0241913), NOAA (NA030AR4600115), and WHOI Academic Programs.

References

- [Balkovsky and Shraiman, 2002] Eugene Balkovsky and Boris I. Shraiman. Olfactory search at high Reynolds number. In *Proceedings National Academy of Science*, volume 99, pages 12589–12593, October 2002.
- [Choset, 2001] Howie Choset. Coverage for robotics - a survey of recent results. *Annals of Mathematics and Artificial Intelligence*, 31:113–126, 2001.
- [Christopoulos and Roumeliotis, 2005a] Vassilios N. Christopoulos and Stergios Roumeliotis. Adaptive sensing for instantaneous gas release parameter estimation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 4450–4456, Barcelona, Spain, April 2005.
- [Christopoulos and Roumeliotis, 2005b] Vassilios N. Christopoulos and Stergios Roumeliotis. Multi robot trajectory generation for single source explosion parameter estimation. In *Proceedings of the IEEE International Conference on Robotics and Automation*, pages 2803–2809, Barcelona, Spain, April 2005.
- [Consi *et al.*, 1994] T. R. Consi, J. Atema, C. A. Goudey, J. Cho, and C. Chrysostomidis. Auv guidance with chemical signals. In *Proceedings of the 1994 Symposium on Autonomous Underwater Vehicle Technology (AUV'94)*, pages 450–455, July 1994.
- [Davies and Gather, 1993] Laurie Davies and Ursula Gather. The identification of multiple outliers. *Journal of the American Statistical Association*, 88:782–792, September 1993.
- [Elfes, 1989] Alberto Elfes. Using Occupancy Grids for mobile robot perception and navigation. *Computer*, 22(6):46–57, June 1989.
- [Farrell *et al.*, 2003] Jay A. Farrell, Shuo Pang, and Wei Li. Plume mapping via hidden markov methods. *IEEE Transactions on System, Man, and Cybernetics*, 33(6):850–863, 2003.
- [Farrell *et al.*, 2005] J. A. Farrell, S. Pang, and W. Li. Chemical plume tracing via an autonomous underwater vehicle. *IEEE Journal of Oceanic Engineering*, 2005. To appear.
- [Ferri *et al.*, 2006] Gabriele Ferri, Emanuele Caselli, Virgilio Mattoli, Alessio Mondini, Barbara Mazzolai, and Paolo Dario. A biologically-inspired algorithm implemented on a new highly flexible multi-agent platform for gas source localization. In *Proceedings of The First IEEE/RAS-EMBS International Conference on Biomedical Robotics and Biomechanics, BioRob 2006*, pages 573–578, February 2006.
- [Ferri *et al.*, 2007] Gabriele Ferri, Michael V. Jakuba, Emanuele Caselli, Virgilio Mattoli, Barbara Mazzolai, Dana R. Yoerger, and Paolo Dario. Localizing multiple gas/odor sources in an indoor environment using bayesian occupancy grid mapping. In *IEEE/RSJ International Conference on Intelligent Robots and Systems*, 2007.
- [German *et al.*, 2008] Christopher R. German, Dana R. Yoerger, Michael Jakuba, Timothy M. Shank, Charles H. Langmuir, and Ko-ichi Nakamura. Hydrothermal exploration with the Autonomous Benthic Explorer. *Deep Sea Research I*, pages 203–219, 2008.
- [Grasso *et al.*, 2000] Frank W. Grasso, Thomas R. Consi, David C. Mountain, and Jelle Atema. Biomimetic robot lobster performs chemo-orientation in turbulence using a pair of spatially separated sensors: progress and challenges. *Robotics and Autonomous Systems*, 30:115–131, 2000.
- [Hayes *et al.*, 2002] Adam T. Hayes, Alcherio Martinoli, and Rodney M. Goodman. Distributed Odor Source Localization. *IEEE Sensors*, June 2002.
- [Ishida *et al.*, 1998] H. Ishida, T. Nakamoto, and T. Moriizumi. Remote sensing of gas/odor source location and concentration distribution using mobile system. *Sensors and Actuators B*, 49:52–57, 1998.
- [Ishida *et al.*, 2001] H. Ishida, T. Nakamoto, T. Moriizumi, T. Kikas, and J. Janata. Plume-tracking robots: a new application of chemical sensors. *Biological Bulletin*, 200(2):222–226, April 2001.
- [Jakuba, 2007] Michael V. Jakuba. *Stochastic Mapping for Chemical Plume Source Localization with Application to Autonomous Hydrothermal Vent Discovery*. PhD thesis, Massachusetts Institute of Technology and Woods Hole Oceanographic Institution Joint Program in Oceanography/Applied Ocean Science and Engineering, February 2007.
- [Kowadlo and Russell, 2008] Gideon Kowadlo and R. Andrew Russell. Robot odor localization: A taxonomy and survey. *The International Journal of*

- Robotics Research*, 27:869–894, August 2008. doi: 10.1177/0278364908095118.
- [Kraetzschmar *et al.*, 2004] Gerhard K. Kraetzschmar, Guillem Pagès Gassull, and Klaus Uhl. Probabilistic quadrees for variable-resolution mapping of large environments. In M. I. Ribeiro and J. Santos Victor, editors, *Proceedings of the 5th IFAC/EURON Symposium on Intelligent Autonomous Vehicles*, Lisbon, Portugal, July 2004. Elsevier Science.
- [Kuwana *et al.*, 1995] Yoshihiko Kuwana, Isao Shimoyama, and Hirofumi Miura. Steering control of a mobile robot using insect antennae. In *Proceedings of the 1995 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 530–535, 1995.
- [Leonard and Durrant-Whyte, 1992] J. J. Leonard and H. F. Durrant-Whyte. *Directed sonar sensing for mobile robot navigation*. Kluwer, Dordrecht, The Netherlands, 1992.
- [Liao and Cowen, 2002] Qian Liao and Edwin A. Cowen. The information content of a scalar plume - a plume tracing perspective. *Environmental Fluid Mechanics*, 2:9–34, 2002.
- [Lilienthal and Duckett, 2003a] A. J. Lilienthal and T. Duckett. Creating gas concentration gridmaps with a mobile robot. In *Proceedings IROS 2003*, 2003.
- [Lilienthal and Duckett, 2003b] A. J. Lilienthal and T. Duckett. Gas source localization by constructing concentration gridmaps with a mobile robot. In *Proceedings EECMR 2003*, 2003.
- [Lilley *et al.*, 1995] Marvin D. Lilley, Richard A. Feely, and John H. Trefry. *Seafloor Hydrothermal Systems: Physical, Chemical, Biological, and Geological Interactions*, chapter Chemical and Biochemical Transformations in Hydrothermal Plumes, pages 357–368. Geophysical Monograph 91. American Geophysical Union, 1995.
- [Lim and Cho, 1992] Jong Hwan Lim and Dong Woo Cho. Physically based sensor modeling for a sonar map in a specular environment. In *Proc. IEEE International Conference on Robotics and Automation*, volume 2, pages 1714–1719, May 1992.
- [Lupton *et al.*, 1985] J. E. Lupton, J. R. Delaney, H. P. Johnson, and M. K. Tivey. Entrainment and vertical transport of deep-ocean water by buoyant hydrothermal plumes. *Nature*, 316(15):621–623, August 1985.
- [Lupton, 1995] John E. Lupton. *Seafloor Hydrothermal Systems: Physical, Chemical, Biological, and Geological Interactions*, chapter Hydrothermal Plumes: Near and Far Field, pages 317–346. Geophysical Monograph 91. American Geophysical Union, 1995.
- [Martin and Moravec, 1996] Martin C. Martin and Hans Moravec. Robot evidence grids. Technical Report CMU-RI-TR-96-06, Robotics Institute, Carnegie Mellon University, Pittsburgh, PA, March 1996.
- [Middleton and Thomson, 1986] J. H. Middleton and R. E. Thomson. Modeling the rise of hydrothermal plumes. Technical Report 69, Canadian Technical Report of Hydrography and Ocean Sciences, 1986.
- [Moravec and Cho, 1989] Hans P. Moravec and Dong Woo Cho. A bayesian method for certainty grids. In *Working notes of AAAI 1998 Spring Symposium on Robot Navigation*, pages 57–60, Stanford, CA, 1989.
- [Moravec and Elfes, 1985] Hans P. Moravec and Alberto Elfes. High resolution maps from wide-angle sonar. In *Proc. IEEE International Conference on Robotics and Automation*, 1985.
- [Nagasawa *et al.*, 1999] Sumito Nagasawa, Ryohei Kanzaki, and Isao Shimoyama. Study of a small mobile robot that uses living insect antennae as pheromone sensors. In *Proceedings of the 1999 IEEE/RSJ International Conference on Intelligent Robots and Systems*, pages 555–560, 1999.
- [Pang and Farrell, 2006] Shuo Pang and Jay A. Farrell. Chemical plume source localization. *IEEE Transactions of Systems, Man, and Cybernetics-Part B: Cybernetics*, 36(5):1068–1080, 2006.
- [Pang, 2004] Shuo Pang. *Reactive Planning and On-line Mapping for Chemical Plume Tracing*. PhD thesis, Department of Electrical Engineering, University of California, Riverside, 2004.
- [Russell, 2001] R. A. Russell. Survey of robotics applications for odor-sensing technology. *The International Journal of Robotics Research*, 20(2):144–162, February 2001.
- [Speer and Rona, 1989] K. G. Speer and P. A. Rona. A model of an atlantic and pacific hydrothermal plume. *Journal of Geophysical Research - Oceans*, 94(C5):6213–6220, May 1989.
- [Thrun *et al.*, 2005] Sebastian Thrun, Wolfram Burgard, and Dieter Fox. *Probabilistic Robotics*. MIT Press, Cambridge, Massachusetts, 2005.
- [Thrun, 2003] Sebastian Thrun. Learning occupancy grid maps with forward sensor models. *Autonomous Robots*, 15:111–127, 2003.
- [Turner, 1986] J. S. Turner. Turbulent entrainment: the development of the entrainment assumption, and its application to geophysical flows. *Journal of Fluid Mechanics*, 173:431–471, 1986.
- [Yoerger *et al.*, 1991] Dana R. Yoerger, Albert M. Bradley, and Barrie B. Walden. The Autonomous

Benthic Explorer (ABE): An AUV Optimized for Deep Seafloor Studies. In *Proceedings of the Seventh International Symposium on Unmanned Untethered Submersible Technology (UUST91)*, pages 60–70, Durham, NH, 1991.

[Zarzhitsky *et al.*, 2004] Dimitri Zarzhitsky, Diana Spears, David Thayer, and William Spears. Agent-based chemical plume tracing using fluid dynamics. In *Lecture Notes in Computer Science*, volume 3228. Springer-Verlag, 2004.