

Planning for AUVs: Dealing with a Continuous Partially-Observable Environment

Dearden, Richard; Saigol, Zeyn; Wyatt, Jeremy; Murton, BJ; Ingrand, F; Rajan, K

Citation for published version (Harvard):

Dearden, R, Saigol, Z, Wyatt, J, Murton, BJ, Ingrand, F & Rajan, K 2007, 'Planning for AUVs: Dealing with a Continuous Partially-Observable Environment', Paper presented at 3rd ICAPS Workshop on Planning and Plan Execution for Real-World Systems, 1/09/07 pp. 87-92. <<http://www.cs.bham.ac.uk/~zas/ICAPSworkshop07.pdf>>

[Link to publication on Research at Birmingham portal](#)

General rights

Unless a licence is specified above, all rights (including copyright and moral rights) in this document are retained by the authors and/or the copyright holders. The express permission of the copyright holder must be obtained for any use of this material other than for purposes permitted by law.

- Users may freely distribute the URL that is used to identify this publication.
- Users may download and/or print one copy of the publication from the University of Birmingham research portal for the purpose of private study or non-commercial research.
- User may use extracts from the document in line with the concept of 'fair dealing' under the Copyright, Designs and Patents Act 1988 (?)
- Users may not further distribute the material nor use it for the purposes of commercial gain.

Where a licence is displayed above, please note the terms and conditions of the licence govern your use of this document.

When citing, please reference the published version.

Take down policy

While the University of Birmingham exercises care and attention in making items available there are rare occasions when an item has been uploaded in error or has been deemed to be commercially or otherwise sensitive.

If you believe that this is the case for this document, please contact UBIRA@lists.bham.ac.uk providing details and we will remove access to the work immediately and investigate.

Planning for AUVs: Dealing with a Continuous Partially-Observable Environment

Richard W Dearden

Zeyn A Saigol

Jeremy L Wyatt

School of Computer Science

University of Birmingham

Edgbaston, Birmingham B15 2TT, UK

{rwd, zas, jlw}@cs.bham.ac.uk

Bramley J Murton

National Oceanography Centre

Waterfront Campus, Empress Dock

Southampton SO14 3ZH, UK

bjm@noc.soton.ac.uk

Abstract

We describe the domain of using an Autonomous Underwater Vehicle to find hydrothermal vents located on the sea floor, and explain some of the difficulties in planning in this domain. We also present a simplified model of the domain, and outline a possible approach for on-line plan generation.

Introduction

The ocean floor was once thought to be lifeless and uninteresting, but since the 1960s, a wealth of interesting geological and biological phenomena have been found. Our interest is in hydrothermal vents—outgassings of superheated water found on oceanic ridges and of interest to both biologists and geologists. However, exploring these depths is difficult, due to the extreme pressures, cold temperatures and lack of light. Instead of using crewed vehicles for exploration, scientific missions increasingly use Remotely Operated Vehicles (ROVs) and Autonomous Underwater Vehicles (AUVs) (Blidberg 2001). These can operate for longer periods than manned craft, and in less accessible areas. AUVs have the additional advantages over ROVs that they do not need a tether to a mother ship, nor 24-hour remote operation. AUVs have proved useful to marine scientists, for example the AUV ABE discovered a series of hydrothermal vents during a research cruise in 2005 (German & Parson 2006).

The key limitation of current AUVs is that they generally either follow a pre-programmed course, or have only rule-based control systems, which can lead to inefficient exploration. We want to develop more advanced planning systems for AUVs, to enable them to maximise the scientific reward of missions, given their uncertain picture of the world, and limited power resources. The task we focus on is to locate and examine hydrothermal vents. The AUV gains clues to vent locations by detecting properties of the outflowing water, which is rich in chemicals, and is of a different temperature and salinity than normal seawater. The AUV will have chemical and temperature/salinity sensors, and from noisy data provided by these it should be able to build a probabilistic map of where vents may be located. The planning aspect of the problem is deciding what the robot should do given the uncertain information it has; for example, should it spend further time exploring a vent it has already examined, move on and try to find further vents, or end its mis-

sion while it is guaranteed to have enough battery power to return to the surface?

In this paper, firstly we describe current techniques used to find hydrothermal vents, and the kind of planning problems that arise when trying to automate this process. We next outline a simplified version of the problem, based in part on our analysis of data from a research cruise where several hydrothermal vents were located, which we hope will be solvable. Finally we describe previous work on similar problems, and suggest an approach for producing a planner capable of finding near-optimal solutions to the (simplified) problem laid out here. We aim to implement such a planner ourselves in the near future.

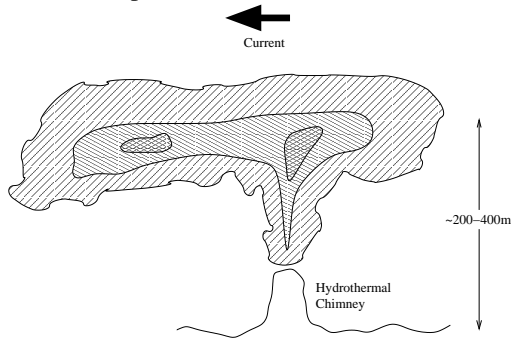
Finding Hydrothermal Vents

Hydrothermal vents are due to plate tectonics, which arise from the layered nature of the Earth. Continental plates are part of the crust, which floats on top of the deformable upper mantle. Vents form on spreading centres, which are plate boundaries where plates are moving apart from one another, and magma flows up from the mantle to form new ocean floor. When cold seawater seeps down through cracks in the rocks on an oceanic ridge, it comes into contact with the hot magma below and heats up. Then convection causes it to rise back up to the sea, dissolving minerals and metals from the surrounding rocks on the way, and it forms a hydrothermal vent when it exits back into the ocean. The most spectacular vents are nicknamed “black smokers” and often exit from tall chimneys formed from minerals they deposit.

Volcanic activity varies along the length of an oceanic ridge, with some areas more active than others, which leads to vents often being found in “vent fields” containing multiple vents in close proximity. Also vents come in different sizes and have different peak temperatures, as some are older than others, and some are formed from water that has not seeped very far down towards the hot magma.

The hydrothermal fluid emanating from vents contains “tracers”, which are unique chemical and physical properties that identify the source of the water, for example the presence of dissolved methane. Fluid rising from a vent will mix with cold sea water, and continue rising and cooling until it reaches a height where it has the same buoyancy as the surrounding water (generally at around 250m). The fluid, still containing tracers from the vent, then disperses

Figure 1: Diagram showing the typical dispersion of a hydrothermal plume (not to scale).



in a plume driven by the current as shown in figure 1. Several issues make it hard to detect this plume: firstly, the vent water is mixed with sea water and is not significantly different in temperature from the surrounding water (however, the interaction between salinity, temperature and density in seawater does mean hydrothermal fluid contains a “temperature anomaly”). Secondly, the current flow is not constant, but will change direction every few hours. Thirdly, sea water often separates into vertical layers with different densities, and hydrothermal fluid will not easily cross these layer boundaries, which means it can sometimes become trapped at a higher or lower altitude than normal. Finally, the water column movement is somewhat turbulent, and tracers will only appear in the plume water intermittently at best. In fact, it is possible for vent tracers to become trapped in a vortex, and there to be no trail leading back to the vent.

Traditional Methods for Locating Vents

This section describes the process by which hydrothermal vents are found, and describes some of the key sensor systems used in locating them. A typical research cruise with the aim of finding vents will use four methods:

- Gathering low-resolution bathymetry data (XYZ data describing the relief of the ocean bottom).
- CTD casts. CTDs are Conductivity/Temperature/Depth sensors, which are attached to a cage containing sampling bottles which is slowly lowered to the ocean floor while keeping the ship on station in a fixed location.
- Sensor platforms towed behind the ship, at a depth of a few hundred metres above the sea floor, comprising CTD, bathymetry, and probably other sensors.
- A submersible such as an ROV or AUV is often used to explore vents close-up.

Bathymetry data is captured by the research ship’s on-board swathe bathymetry sensors, which can map a region several kilometres wide in a single pass. Bathymetry can be used by experts on vent geology to identify likely locations for vents based on the relief of the area, and can guide deployments of the CTD and deep-tow platforms. CTD and deep-tow data will hopefully locate vent plume(s), which provide pointers

to where vent fields are. Vents are then pinpointed by low-altitude surveys in a submersible.

CTDs, deep-towed platforms and ROVs/AUVs can all make use of similar chemical sensors to track down vents. The important sensors are (Baker, German, & Elderfield 1995):

Optical backscattering When the hydrothermal fluid meets cold sea water, many of the minerals precipitate out into particle form, making the vent water very cloudy. The particle concentration in the water can be estimated by measuring the amount of light reflected back using a Light Scattering Sensor (LSS, or Nephelometer). The background level is nearly constant below 1000m, which makes the readings easy to interpret.

Methane Dissolved methane is an unambiguous indicator of vent activity and is easily measured.

Reduction potential (often referred to as Eh). Redox potential measures the chemical reactivity of the water; hydrothermal water has low redox potential, as it is low in oxygen. Eh is useful as strong changes in it are only observed close to a vent, within a few hundred metres.

Manganese, Iron, and other metals Manganese in particular is enriched by several orders of magnitude in hydrothermal fluid compared to normal seawater.

Potential temperature versus salinity anomalies While the temperature of hydrothermal fluid far from the vent is indistinguishable from that of the surrounding water, it will be identifiable by a change in the normally linear relationship between potential temperature and salinity in a given region of ocean. Such anomalies can be found by examining a series of CTD data, as salinity and potential temperature are just functions of conductivity, temperature and density.

Smart AUVs

Existing AUVs are extremely useful tools for searching for hydrothermal vents, and contain intelligent software to allow them to avoid obstacles, navigate a fixed route regardless of current direction and magnitude, and in some cases to hover in a fixed location. However, they are generally not capable of deciding on a mission route for themselves, or if they are, use only a basic algorithm for exploring previously detected areas of high tracer activity. Our objective is to replicate the overall behaviour of the traditional approach for finding vents described above, beginning once a promising area of seabed has already been identified. The scenario we envisage is that a research vessel will cruise over the area being explored and identify likely areas for vents. These areas will then be explored by deploying an AUV. The AUV will attempt to locate vents and perform science until its battery is exhausted, at which point it will return to the surface.

AUVs for such a task should decide when to use an exhaustive search pattern, when to investigate an area of high vent probability, and if investigating potential vent areas, which areas it would be most fruitful to explore further. They should decide when to survey at high altitudes, looking for far-off vents, and when to search for near vents at

low altitudes. They should decide when a vent has been localised to a high enough probability, and examined thoroughly enough, such that more reward would derive from searching for a new vent. They should decide when the risks of not having enough resources to return to the surface outweigh the potential science benefits of continuing the mission. They should decide when to take science actions such as taking a photograph or collecting a water sample.

Working Version of the Problem

One way to formulate the planning problem is as a partially observable Markov decision process (POMDP). A POMDP is a tuple $\langle S, A, O, T, H, R \rangle$ where S is the set of possible system states, A is the set of possible actions, O is the set of possible observations, $T = P(s, a, s')$ is the transition function that governs how an action changes the state of the system, $H = P(o|s, a)$ is the observation function that governs how likely each observation is given a state and action, and $R = r(s, a)$ is the reward function which specifies the immediate utility of doing a particular action in a state.

In our domain, the state space S consists not only of the state of the vehicle, but also the (unknown) locations of all objects of interest. Unlike typical POMDPs, this is a mixed discrete-continuous space as the vehicle could be at any location, as could an unknown number of vents. In addition, we will want to reason about how much of our resources (battery power and time) are available at any stage, so we may need to include these in the state as well.

A is the set of actions the vehicle can perform, consisting of both movement actions and sensory actions such as measuring the temperature of the surrounding water. Again, there is a continuous space of possible actions that could be performed. As an additional complication we may wish to account for the time taken performing an action as part of our decision-making. For simplicity we will assume the vehicle has access only to two sensors, but as each of these produces a continuous measurement, the observation space O is again continuous.

Depending on exactly how we formulate the planning problem, the transition function T can be thought of as only applying to movement actions as all the others don't change the state, they only provide information. However, if we include resource variables into our model in the same way as (Bresina *et al.* 2002) do, then even information-gathering actions have an effect in terms of taking time and using battery power, so they too have an effect on system state. Because of the uncertainty produced by currents, the transition function for movement actions will tend to be quite noisy. Resource usage may also be quite uncertain.

The observation function H for movement actions contains very little noise for many underwater vehicles as the Doppler sonar used to measure position is very accurate. However, as we will see below, the observation function for the other instruments can be extremely noisy. Even given a set of perfect sensor readings, the set of potential source vents for these readings is large, which will cause substantial problems in interpreting the observations.

Finally we consider what the reward function R should look like for this problem. The overall goal of a mission of

this kind is to return science data to the surface. That is, collecting the data (finding vents and doing science on them) has no intrinsic value, only returning that data has value. In addition, there are serious risks involved with some activities. The temperature of the water leaving the vents is enough to melt parts of the vehicle, so there needs to be a significant negative reward for vehicle damage.

Plume Model

We have developed a simplified model of the dispersion of tracers from a hydrothermal vent. This model has been generated from two main sources: our intuition about the behaviour of plumes, and the analysis of survey data from a research cruise. The data is from RSS Charles Darwin cruise CD169, which managed to locate a hydrothermal vent on the Mid-Atlantic Ridge during February and March 2005 (German & Parson 2006). The data was acquired from three separate instruments: firstly a deep-towed side-scan sonar platform (TOBI), secondly CTD (Conductivity/Temperature/ Depth) data from 15 locations, and finally from sensors on an AUV, ABE, which is operated by the Woods Hole Oceanographic Institution (WHOI) and was sent on 6 missions from the Charles Darwin.

We chose to model just two variables initially, optical backscattering (LSS) and reduction potential (Eh), as these are two of the most important indicators of hydrothermal activity, and were available from most of the data sources. Figure 2 shows these two variables plotted with distance from a vent site, from ABE data recordings at an altitude of 5m.

Note the noise for LSS and Eh is very different—for LSS, the signal is approximately either present or missing, suggesting that the particles from the hydrothermal fluid have not mixed uniformly with seawater. Redox potential, however, shows a much smoother variation.

Our model deals with two separate situations, depending on whether or not the AUV is directly above a vent. In each case, three distributions are used to generate an "expected" value for the sensor, given the location of the AUV relative to the vent, and the AUV's estimate of the current (which it can calculate from its path over the ground and its motor settings). These are a horizontal distribution in the direction of the current, a horizontal distribution perpendicular to the current, and a vertical distribution. Expected tracer values are found by intersecting all three curves.

The distributions along the current direction above a vent were found by examining LSS and Eh data from an altitude of 5m above the ocean floor, and fitting them to a sum-of-Gaussians based model. The resulting fit is shown in figure 2. The vertical distribution was found from looking at LSS data from a CTD cast approximately directly over a vent, and again fitting it to a sum-of-Gaussians model, as shown in figure 3.

The remaining four distributions needed to model the plume are shown for LSS in figure 4. Figure 4 (a) shows an approximation of the behaviour LSS directly above a vent, perpendicular to the current. The other subfigures describe the model away from a vent, i.e. outside the 50m by 10m area defined by the limits of the distributions in figures 2 and 4 (a). The vertical profile is shown in figure 4 (b), and

Figure 2: Distributions used to approximate backscattering and Eh dispersal over distance, at 5m altitude. We believe that the sharp peak on the left represents a separate vent located somewhat off the track plotted. Eh values have been negated for ease of visualisation, so the peak shown at $-2.2Eh$ is in fact a trough at $2.2Eh$.





Figure 3: Attenuation of the LSS signal with altitude directly above a vent



represents the general form observed in CTD casts which were not directly above a vent. This distribution is essentially a cross-section of the buoyant plume as shown in figure 1, away from the "tail" of the mushroom cloud.

The horizontal distributions away from a vent are shown in figure 4 (c) and 4 (d). These are based on estimates of the range of tracers as they are dispersed horizontally in the plume, and are again modelled as sums of Gaussians. The distributions used for redox potential will have the same form as those shown for LSS.

The model presented above will frequently not match measurements from a real vent, because hydrothermal vents are very variable in terms of the characteristics of the water emitted from them. This variation is found within a vent field, in the peak values of each tracer, the geographic spread of tracers, height of the buoyant plumes, and also between vent fields and especially between fields found in different oceans. For example, in the Pacific ocean, the buoyant

plumes are slightly warmer than surrounding water, whereas in the Atlantic they are slightly colder (Baker, German, & Elderfield 1995). Despite this, our model is a sensible first step, and we hope that the addition of noise when formulating our observation function H will make it reasonably robust to the diversity of real-world vents.

Previous Planning Work

There are a number of relevant approaches in the planning literature. If we look at the POMDP literature, algorithms such as incremental pruning (Cassandra, Littman, & Zhang 1997) and bounded policy iteration-based approaches (Poupart & Boutilier 2004) have been used to solve POMDPs with tens of thousands of states. However, few of these address the problems of continuous state and action spaces, and continuous resources.


On the other hand, the Mars rover work of (Bresina *et al.* 2002; Dearden *et al.* 2003; Mausam *et al.* 2005) tackles the problem of planning with continuous state variables for resources such as battery power, and uncertainty over resource usage. However these algorithms work for the completely observable case.

Frequently in planning we deal with continuous state spaces and uncertainty by abstracting them away, and this is certainly an approach that could be taken here. If we look at how human-controlled searches for vents are performed, they tend to conform to a relatively simple pattern where areas are systematically searched by 'mowing the lawn' at an altitude of 3-500 metres off the seabed to produce a large-scale map of plume density, following which likely vent locations are determined and lower ($\sim 5m$ altitude) searches are performed to localise the vents. Simple current-following algorithms may also be used to follow a signal to its source. Each of these activities can be thought of as an abstract action and a planner used to decide between them. They can then be decomposed into sequences of lower-level actions which can be interrupted when significant observations are made. This approach is closely related to HTN planning approaches (Erol, Hendler, & Nau 1994). The advantage is that the planning task becomes relatively simple. However, the reason human-controlled missions proceed in this fashion is that data tends to be analysed off-line after it has been collected, so a search consists of a high-altitude mission, followed by analysis of the data collected to identify promising regions, which are then searched at low altitude in a subsequent mission. Because we are doing the analysis on-board while data collection is continuing, we should be able to do significantly better than these approaches.

Our Planning Approach

The overall problem can be subdivided into the challenges of state estimation and planning. In practice these two should interact with each other more than just the obvious interaction that the estimated state is an input to the planner. Since we are exploring, changing the state estimate is a significant goal of the planner and the planner needs to reason about the belief state. In addition, there are significant differences between our current model and the actual behaviour of the

Figure 4: Typical LSS distributions



(a) Variation perpendicular to the current, directly above the vent (5m altitude)

(b) Variation with altitude away from the vent (cross-section of the buoyant plume)

(c) Variation in the current direction at the altitude of the buoyant plume

(d) Variation perpendicular to the current at the altitude of the buoyant plume


plume over a vent or vent field. Separating the state estimation component from the planner allows the planner to be independent of the model, and therefore we can leave the planner unchanged when we improve the fidelity of the observation model. For now, we will set these arguments aside and consider the two systems independently.

State Estimation

The challenging part of the state estimation problem is estimating vent locations. We get a stream of observations from the vehicle's suite of instruments and must build an estimate of where we are likely to find vents. Essentially this is a mapping problem similar to that in SLAM (Thrun 2006), but here the possible explanations for an observation aren't independent as they are in SLAM. The model we developed above to predict the plume characteristics from a vent is too complex to invert to predict vent locations from observations so we propose a Monte Carlo approach to state estimation. Consider a search where we know there is exactly one vent. We can sample a number of possible vent locations from our prior distribution and compute for each of them what the characteristics of the plume might be, using our observation model in the forward direction. As we collect observations with the vehicle, samples that predict the actual observations poorly are discarded, and those that predict well are multiplied by resampling. The result is a set of sampled vent locations that with appropriate action choices should converge on the true vent location. To cope with the fact that we have a continuous state space, we smear out the well performing samples over nearby locations by transforming the set of point samples into a continuous distribution, for example by treating each sample as the mean of a weighted Gaussian, and then sample from the Gaussian mixture.

In practice there could be multiple vents all contributing to the plume, so the model proposed above is unrealistic. The possibility of multiple vents means that the state space consists of the locations of an unknown number of vents. As the number of vents rises, the dimensionality of the state space makes it impractical to sample from. We propose to solve this problem using a clustering approach similar to the factored particle filtering approach of (Ng, Peshkin, & Pfaffer 2002), but with the clusters being dynamically recom-

Figure 5: (a) The result of state estimation: a map of the likelihood of a vent at any location. (b) A PRM generated from the vent location map.



puted as the vehicle moves. For reasons of space we omit the details here, but the final result is, for every possible event location, an estimate of the probability of a vent existing at that location. Figure 5 (a) shows an example density.

Criteria

Imagine that we have constructed a probability map of the type described above. We cannot allocate values directly to the physical states (vehicle poses), because reward must be a function of information, and we are planning the gathering of future information. There are several possibilities related to the value of increased confidence about vent presence and location within an area. We might reward the reduction in the sum of the entropy at each point in the map. This is similar to the infotaxis approach of Vergassola et al. (Vergassola, Villermaux, & Shraiman 2007) where the objective is to maximise the local reduction in the entropy of the belief density. In our case we have to reason about the entropy of a simpler distribution but at every point in the world.

Another criterion, and one that would induce rather different behaviour when optimised, would be to minimise the probability of failing to spot vents (minimise the number of false negatives) within a specified search area. If optimised for certain search field shapes this could naturally generate often observed strategies such as spiral search, or mowing the lawn. Finally, whichever of these criteria we use, we also need to reward for the likely value of the observations taken from each vent. Suppose, for example, that the vehicle is rewarded for pictures of a vent, with decreasing rewards for

increasingly similar views to those already taken. Those rewards must also take into account the likelihood of the presence or absence of the vent in the view, thus including the effect of the likelihood of taking a picture of a non-vent.

Planning

As we saw previously, we can represent this problem as a POMDP over a continuous state, action, and observation space. Since even discrete POMDPs are intractable for large state spaces, it's clear that we can't hope to solve such a POMDP optimally. Fortunately, we are generally interested in a satisficing rather than optimal solution as the intention is to trade off coverage for quality—the AUV can't compete with an ROV at the exploration and science, but since it doesn't require the continuous presence of a research ship, and multiple vehicles can be deployed, the amount of seabed that can be explored is much greater.

The question is how best to produce a satisficing solution. Here we examine two contrasting approaches. In the first, we do a minimal amount of approximation in an effort to build a system that continually integrates new observations and performs close to optimally but that may potentially behave very differently than a human-planned mission might. In the second we take a more abstract approach where human-designed behaviours become the actions we are planning with. It is possible that the best approach—taking into account both plan quality and the computation required to reach it—lies between these two extremes.

Probabilistic road maps (PRM) (Latombe 1991) are a common approach to the problem of path planning in continuous spaces. The idea is to sample a random graph in the space and only allow movement over the edges in the graph. The approach is typically used in obstacle avoidance, but can easily be adapted to our domain. Since our state estimate provides a distribution over likely vent locations, and we want to spend much of our time in the highest probability locations, when we generate the vertices in the PRM graph it makes sense to place more of them in these high probability areas, so we bias the PRM graph by sampling from the vent location probability densities. An example PRM graph is shown in Figure 5 (b).

The PRM graph results in a small discrete number of movement actions and possible vehicle locations rather than the continuous spaces in the original problem. Since the edges in the graph may be of different lengths, each will have a different cost in terms of resources used. We label the edges with their expected resource usage and take this into account when selecting which action to perform. If the vehicle has significant costs to change direction, the vehicle pose can be included in the PRM graph as an extra dimension attached to each vertex, and the edge costs will be appropriately higher if direction changes are required to follow them. Similarly, water currents may affect resource usage and must be taken into account in edge costs.

Another issue is the continuous resource space. The approach taken in (Dearden *et al.* 2003) is relevant here—resource usage is regressed through actions to estimate the set of goals achievable given the current state and resources. Unlike their Mars rover domain, we don't know with cer-

tainty where the 'goals' or high-reward states are, so it is harder to do the analysis. The resource usage distributions will be much less uncertain here.

Conclusions

In this discussion paper we have described a new and challenging real-world domain for planning, discussed its formalisation, and sketched the types of algorithms that we believe might be applicable. The key aspects of the problem are that the goal is to collect information, while the state, action and observation spaces are continuous and multidimensional. While we have presented an underwater application, the same ideas are clearly applicable to other autonomous science domains, such as planetary rovers.

References

- Baker, E. T.; German, C. R.; and Elderfield, H. 1995. Hydrothermal plumes over spreading-center axes: Global distributions and geological inferences. In Humphris, S.; Zierenberg, R.; Mullineaux, L.; and Thomson, R., eds., *Seafloor Hydrothermal Systems: Physical, Chemical, Biological, and Geological Interactions*, number 91 in Geophysical Monograph. American Geophysical Union. 47–71.
- Blidberg, D. 2001. The development of Autonomous Underwater Vehicles (AUV); A brief summary. *Proceedings of ICRA 2001*.
- Bresina, J.; Dearden, R.; Meuleau, N.; Smith, S. R. D.; and Washington, R. 2002. Planning under continuous time and resource uncertainty: A challenge for AI. In *Proceedings of UAI-02*, 77–84. Morgan Kaufmann.
- Cassandra, A.; Littman, M. L.; and Zhang, N. L. 1997. Incremental Pruning: A simple, fast, exact method for partially observable Markov decision processes. In Geiger, D., and Shenoy, P. P., eds., *Proceedings of UAI-97*, 54–61. Morgan Kaufmann.
- Dearden, R.; Meuleau, N.; Ramakrishnan, S.; Smith, D. E.; and Washington, R. 2003. Incremental contingency planning. In *Proceedings of ICAPS'03 Workshop on Planning under Uncertainty and Incomplete Information*.
- Erol, K.; Hendler, J. A.; and Nau, D. S. 1994. UMCP: A sound and complete procedure for hierarchical task-network planning. In *Artificial Intelligence Planning Systems*, 249–254.
- German, C., and Parson, L. 2006. RRS "Charles Darwin" cruise 169, 17 Feb-19 Mar 2005. Hydrothermal exploration of the southern Mid-Atlantic Ridge. NOCS Cruise Report 6, Southampton.
- Latombe, J.-C. 1991. *Robot Motion Planning*. Kluwer International Series in Engineering and Computer Science. Kluwer Academic.
- Mausam; Benazera, E.; Brafman, R.; Meuleau, N.; and Hansen, E. 2005. Planning with continuous resources in stochastic domains. *Proceedings of IJCAI 05* 1244–1251.
- Ng, B.; Peshkin, L.; and Pfeffer, A. 2002. Factored particles for scalable monitoring. In *Proceedings of UAI-02*.
- Poupart, P., and Boutilier, C. 2004. VDCBPI: an approximate scalable algorithm for large scale POMDPs. In *Advances in NIPS 17*, 1081–1088.
- Thrun, S. 2006. Simultaneous localization and mapping. In Jefferies, M., and Yeap, W.-K., eds., *Spatial Mapping Approaches in Robotic and Natural Mapping Systems*. Springer.
- Vergassola, M.; Villermaux, E.; and Shraiman, B. I. 2007. 'Infotaxis' as a strategy for searching without gradients. *Nature* 445(7126):406–409.