Planning for Autonomous Underwater Vehicles

Zeyn Saigol

Intelligent Robotics Lab meeting

July 31, 2007

Outline

Outline

Oceanography Aims

- Hydrothermal vents (Black Smokers) - discovered in 1977
- Plate tectonics - 1960s
- Mechanism for most interesting ocean behaviour
- Vents are formed at spreading centres such as the Mid-Atlantic Ridge

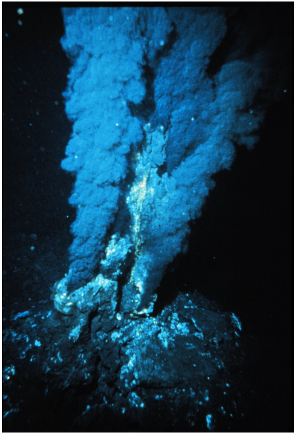

Autonomous Underwater Vehicles

- Unmanned, untethered submersibles
- Autosub, developed in Southampton

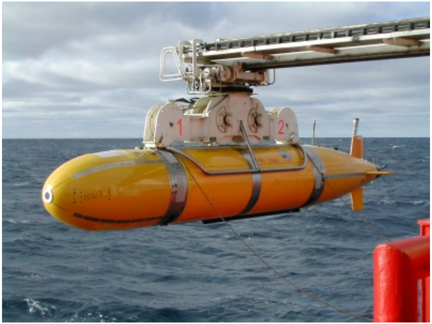

- Cheaper than manned vehicles
- Can get to places tethered vehicles can't
- No need for human supervision - extended missions
- Current control systems are simple if-then-else rules

Outline

Planning Aims

- Richard's work on Mars rovers
- Key characteristics of problem
  - Continuous state space
  - Uncertain action outcomes
  - Finite-duration actions
  - Limited resources for executing the plan
- Additionally, in the AUV problem, we have a partially- observable state space.

Outline

Oceanography Data

- Collaboration with Bramley Murton @ National Oceanography Centre, Southampton
- Data from research cruise on Mid-Atlantic Ridge
- Sources include:
  - Vertical profiles from dropping a sensor rosette over the side of the ship
  - Readings from a sensor platform towed behind the ship, a few hundred metres above the sea floor
  - Sensor logs from an AUV
- 1.5 Gb of data in total

Sample of NOCS Data

- Example of data from the AUV, recorded at 5m altitude

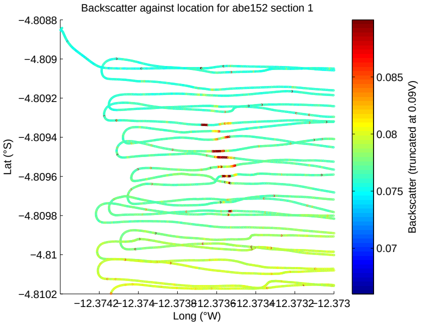

- Data is for backscattering of light - indicates density of particles in the water

Outline

Definition as an MDP

- Problem can be defined as a Markov Decision Process
- State S contains
  - Pose of AUV
  - Battery power remaining
  - Locations of nearby hydrothermal vents
- Actions A are to move the AUV to a nearby waypoint
- Transition function $T = P(s, a, s')$ models the motion of the vehicle and battery use. With waypoints, there is little uncertainty in the new location, but lots in the power used
- Reward function R consists of
  - Small positive reward for finding each unique vent
  - Large negative reward for exhausting battery power before returning to the surface

Solving MDPs

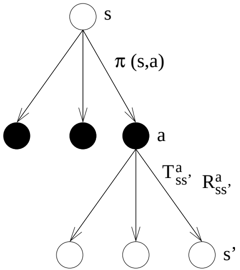

- $\pi(s, a)$ - policy of agent: probability of choosing action a from state s
- $T^{a}_{ss'}$ - transition function
- $R^{a}_{ss'}$ - reward function

- Bellman Equation

$$V(s) = \sum_{a} \pi(s, a) \sum_{s'} P^{a}_{ss'} \left[R^{a}_{ss'} + \gamma V(s') \right]$$

Outline

Unknown State - POMDPs

- We can't observe the location of vents directly => in fact have a Partially-Observable MDP (POMDP)
- We make observations *o* of the temperature, optical backscattering, etc
- Can use an Observation Model $H = P(o|s)$ to generate a probability distribution of vents over the region

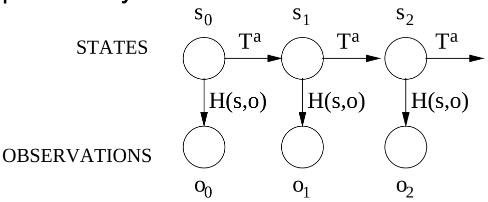

Observation Model

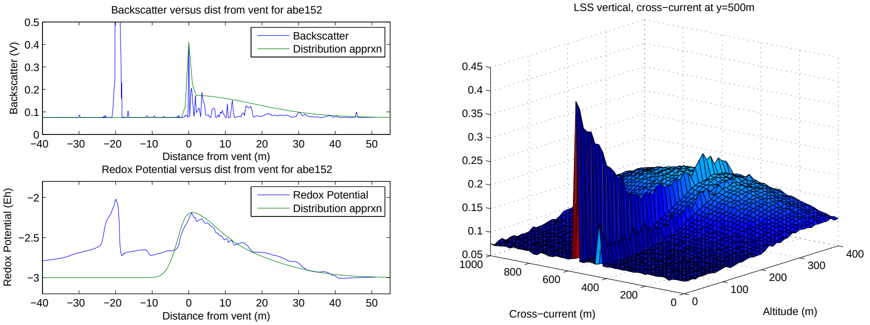

- Mixture-of-Gaussians model
- Create observation function by adding Gaussian noise

POMDP Approach

- Use Bayes' to invert the observation model and generate a belief state $b(s)$, probability of being in true state s
- Can then use this to reduce the POMDP to an MDP in "belief state space" (which is continuous even for a discrete state space)
- Find that the optimal action depends only on the belief state, not the true state
- The problem splits into two parts:
  - State estimation (calculating the belief state)
  - Planning

Outline

State Estimation

- Can accurately track AUV pose (using doppler sonar) and battery life
- Issue is estimating the location of nearby vents - very uncertain
- The belief state can be found from the Bayes filtering equation

$$b'(s') = \alpha H(s', o) \sum_{s} T(s, a, s') b(s)$$

- If we had a linear observation function with Gaussian noise, this could be solved exactly using a Kalman Filter

Particle Filters

- Approximation technique for finding the belief state for arbitrary observation models and noise
- Idea is to maintain a set of samples of the state, and use the Bayes filtering equation to update these samples over time

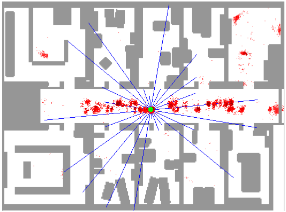

- At each time step, particles are moved according to the transition function, and then re-sampled according to their likelihood under the observation function

Independence Issue

- Unfortunately, we cannot use a straight Particle Filter
- This is because we will probably have multiple vents, and each vent con tributes to the sensor readings we see

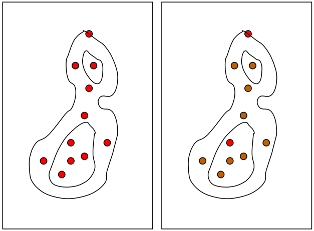

- This means we *cannot* treat the state space samples as *independent*
- Probably need a PF-based solution factored by number of vents

Outline

Biologically-Inspired Algorithms

- Given a prevelant current/wind direction, the problem of localising the source of a chemical plume is solved frequently in nature
    - Male moths follow an intermittent pheromone signal to locate mates, using an algorithm that has been proven to be very effective in real-world experiments

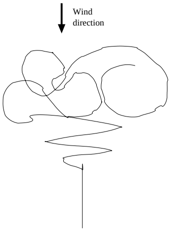

- AUV situation is more complex because have multiple sensors, resource constraints, etc

PRM Path-Planning

- The key part of the planning problem is to find vents
- As such, it is a path-planning problem, and a PRM algorithm is applicable

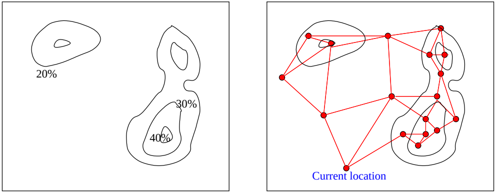

- Our idea is to weight the state-space samples by the probability of a vent being present at that point

Summary

- AUVs pose an interesting planning challenge, due to partially-observable continuous state space
- The problem of finding interesting oceanographic phenomena decomposes into a state estimation problem and a planning problem
- Both of these are hard!

- Future work
  - First task - understand state estimation better
  - Use Moravec's cell occupancy methods to examine in detail
  - Possibly switch to a batch-learning approach