

Planning and implementing trajectories for auv's to track evolving ocean processes

Where is this happening?

Southern California -- want to monitor the ocean near LA

First few pages - general info on auv's

Uses the ROMS - regional ocean modeling system

Think of it as a black box from which we can get information about how points in the ocean are moving

2-way communication -> give data to the black box which will improve it's prediction, helps in the long term, but not in the short term

don't know how the data is used to improve this model

Platform: AUV glider

can't steer underway, can only "aim and shoot"

consider to have visited a waypoint once we've shot at it

because it uses average current, error is often large (1.1km on 2km trajectories)

basically, because you'll almost never make it to your ideal waypoint, you just go ahead and aim for the next waypoint (but don't assume you surfaced where you thought you would) so error doesn't accumulate.

does all it's turning in the first yoyo, then moves in a planar path

dive is completely blind - could get pushed sideways

if there's no current, you could go any direction

could we measure how we're getting off course? No, because the IMU is not accurate enough -- too much noise

just short of a passive drifter

but can stay in water for weeks at a time

sacrifice positional accuracy

coded like it can be commanded precisely but this is not the case

Algorithm:

general idea: aim for the edge/center in order to track it

Motivation

track sudden growths of harmful algae - HABs

although process is 3D, for simplicity consider only 2D extent

consider only one convex region of interest (convex hull, to have centroid)

assume centroid of convex hull is centroid of area of interest

Waypoint selection algorithm

Centroid tracking

find the where the centroid will be when you surface, aim for there

default distance - how far the glider can travel in 4 hours, have upper and lower bounds on distance close to this distance

3 conditions: if the centroid will be within the range, just aim for it

if it's closer, take a detour to an additional waypoint within the ellipse which won't make you go over the desired distance -- basically we can explore more and still reach the centroid, so might as well

there's not really any justification for this, but might as well gain more information

problem: we know something about how the currents work, but we're doing calculations as if it took the same amount of time to go in any direction.

the thing is, if we incorporated that information, it would become difficult to calculate.

if it's farther, aim for an even more future centroid, and head in that direction as far as you can

Boundary tracking

try to find drifters where you can travel the ideal distance to reach them (could have one or two)

if none, find the average angle towards the area of interest, travel that direction

Combining these two

put drifters on the plume, plan using the algorithms, go!

Results

not what a human operator would choose, but the paths allow the gliders to stay in the area of interest

but, no good metric for what constitutes an optimal path

in the simulation, this looks good, because the model used to calculate where the boat should aim is the same as that used to figure out where it comes out

note: the ocean model is not actually that great

in real situation, given some waypoints, got pretty good accuracy

note, this is in a relatively quiet area

note: exists a model which reduces error

Challenges

velocity of current changes with depth, makes prediction difficult

hard to figure out when the HAB starts and ends

models for ocean behavior are not accurate

RESL is trying to convince people that these could be real robots

Can't dive again until above 25m deep, need at least a few meters to start diving properly

Use of fins? Might help, but might take too much energy

The users don't really care about 1km error

There are underwater vehicles that are powered which do more sophisticated location

Centroid tracking: is this the thing to do, or is there something else?

something being done: use lagrangian drifters, use data about where they went to drive other AUV's following them, could be more useful

Criticism, why not just drop passive drifters into a plume, if you know where it is?

no coherence, don't always know where it is

Criticism, why not do planning which incorporates uncertainty in the plan

no useful measure of uncertainty

some other work to learn an uncertainty model on top of ocean model

uncertainty is a problem, could use information about it to do long-term planning

to get good certainty later at the cost of uncertainty now

next paper - look at uncertainty and waypoint generation

Inspection planning for sensor coverage of 3D marine structures

Goal

achieve complete sensor coverage of a complex structure using an agent with many degrees of freedom

motivating situation: ship hull inspection with a bathymetric sensor

plan on a triangular mesh model

Related work

art gallery problem - theoretical framework

car painting/ building inspection - geometry

challenging issues in computational geometry for complex 3D models in the absence of a priori knowledge

from the art gallery p.o.v. gain information only at nodes

this system uses the postman framework, get information along paths as well

Set cover + traveling salesman

Sensor footprint

Maximum range and angle, can calculate which points are visible along an edge

Algorithm

Initially, do similar to probabilistic roadmaps, random views

Then, once unobserved size is small enough, identify locations to identify missing views

For initial step

- find feasible nearby views, remove from an unobserved set, if no collisions

For add to graph

- find a random sample, find its feasible neighbors, add to graph

Add missing view,

- get an unobserved point

- until you find a view which can see that point, draw a random point and do a

Newtonian search to try to find a view which works

- find an edge which observes that view, add it to the graph in a feasible location

But Newtonian computation is quite expensive, so want to do the random search first

So given this dense graph, need a walk in the graph allowing you to observe the entire thing without visiting every edge

- consider as optimization

- minimize cost of moving along edges, subject to 1. flow constraints 2. All nodes inspected 3. Initial node is on graph 4. Integer

Do subtour connection in case there are multiple tours

- take out the one containing the initial position

- use dijkstra to find nearby nodes on tours and connect together

- need to do it this way rather than subtour elimination. otherwise, would have to add exponential constraints. No performance guarantees

Results

- every node in the structure must be viewed

- get decent paths, computation time seems high...

Criticisms

- dynamics constraints?

- have to keep moving while doing things

- for static inspection, no hurry in most cases, might as well do it slowly

- tracking an edge might not work fully

- suggest doing a dataset with moving the object to compare to moving the robot

ship has much higher difference between 95% and 100% -- more complex, trickier to observe

- interesting to see how badly a quick algorithm would do

- majority of time is generating initial covering graph

- optimize missing view -- save some time

- manage the division of time between graph construction and path-finding

Structure size matters...

Uniform coverage not important here, but uniform information about the surface is!--
does resolution matter?