Research Statement – Planning With Uncertainty

Nicholas Roy
Supervisors: S. Thrun & T. Mitchell

Problem: One essential component of any operational mobile robot system is the ability for the robot to localize itself; that is, to determine its position in space consistently and accurately from sensor data. There are many successful localization methods that are able to determine a robot's position relative to a map using sonar, laser and camera data.

However, most localization methods fail under common environmental conditions. For example, localization error can arise when a mobile robot follows a path through a wide-open space, such as outdoors, or in a very large or crowded room. Cameras can also fail in regions which lack sufficient visual structure, such as blank walls or ceilings. Since these environmental conditions are relatively common, a mobile robot navigating reliably in the real world must allow for the potential failure of its localization methods.

Impact: The motivation for coastal navigation is generating trajectories for the mobile robot that reduce the likelihood of localization error. Our solution is to have the robot identify how good regions of space are for localization, and alter the path planning accordingly. We have developed a planner that uses maps of the environment, where each cell in the map contains a notion of information content available at that point in the map, that corresponds to the ability of the robot to localize itself. The higher the information content, the better the localization ability.

State of the Art: Developing motion planning algorithms based on positional uncertainty is not a new idea. Erdmann developed such strategies as well as probabilistic strategies. Work has been done on trajectory generation with respect to positional uncertainty. Takeda et al. [Takeda et al., 1994] do not use the localization process to generate the positional uncertainty across space, but generate probability distributions based on an explicit model of the sensor, and therefore is not generalizable across sensors. Furthermore, the environment is assumed to be static, so the effect of dynamic obstacles on localization is not modelled.

Considerable work in the field of partially observable Markov decision processes (POMDPs) [Koenig and Simmons, 1996] has allowed many mobile robots to model positional uncertainty explicitly. The POMDP process is in fact mathematically elegant solution; a POMDP produces a policy for all states in the environment and can therefore be used to produce the optimal policy for any desired goal location and uncertainty state, assuming such a policy exists. However, one drawback to the use of traditional POMDPs is that the computation time is exponential in the number of states. In a real-world environment, POMDPs are completely computationally intractable. Essentially our solution operates by reducing the high-dimensionality of the POMDP to two or three dimensions, allowing the use of a dynamic programming approach to find the policy for the low-dimensioned space.

Approach: Our solution to these problems is inspired by traditional navigation of ships. Ships often use the coasts of continents for navigation in the absence of better tools such as GPS, since being close to the land allows sailors to determine with high accuracy where they are. The success of this method results from coast lines containing enough information in their structure for accurate localization. By navigating sufficiently close to areas of the map that have high information content, the likelihood of getting lost can be minimized.

We have developed a measure of the information content of the environment based on the results of the localization process. Let the robot's position be given by the initial probability distribution, $P_{\mathbf{x}}$, defined over $\mathbf{X} = (X, Y, \Theta)$. The robot acquires a sensor measurement $\mathbf{s}$, for example a set of range data from a laser sensor. The localization process takes $P_{\mathbf{X}}$ and the sensor data, $\mathbf{s}$, and returns the posterior probability distribution $P_{\mathbf{X}|\mathbf{s}} = p(\mathbf{x}|\,\mathbf{s})$, again defined over the space of poses of the robot, (X, Y, Θ).

The entropy, $H(P_{\mathbf{X}})$, of a probability function, $P_{\mathbf{X}}$, provides a good measure of the certainty with which the robot is localized. The entropy of a probability distribution, $P_{\mathbf{X}}$, is computed over the space of all possible poses (X, Y, Θ) and is defined as:

$$H(P_{\mathbf{X}}) = -\int_{\mathbf{X}} p(\mathbf{x}) \log(p(\mathbf{x}))\, d\mathbf{x} \tag{1}$$

For a particular pose distribution, we can compute the information content, I, of the robot's current position by computing the difference between the expected entropy of the positional probability conditioned on the sensor

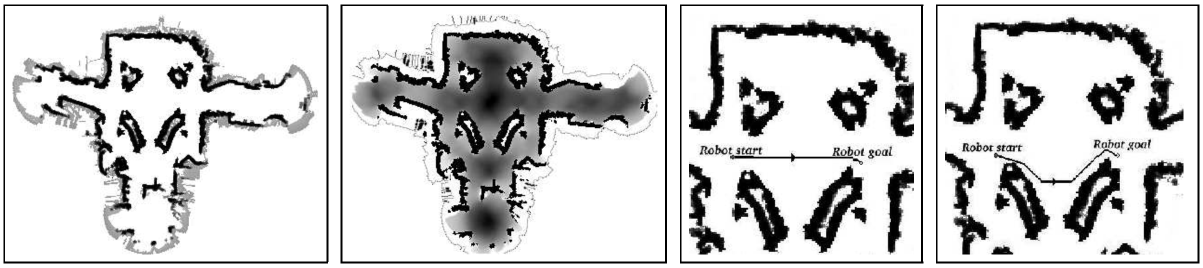

(a) Museum Map (b) Entropy Map (c) Conventional Trajectory (d) Coastal Trajectory

Figure 1: On the left is the map of the National Museum of American History, and its associated entropy map. Example trajectories using the conventional and coastal planners are shown on the right. Note the motion of the robot along the wall for the coastal planner.

measurement, $E(H(P_{\mathbf{X}|\mathbf{s}}))$ and the entropy of the prior position distribution, $H(P_{\mathbf{X}})$:

$$I = E(H(P_{\mathbf{X}|\mathbf{S}})) - H(P_{\mathbf{X}}) \tag{2}$$

We have ignored the issue of the prior probability distribution of the robot's position. The entropy computation is heavily dependent on the robot's prior belief in its position, $p(\mathbf{x})$. Modeling robot navigation as a partially observable Markov decision process, or POMDP, would be one method for handling this dependency [Kaelbling et al., 1998]. However, the POMDP requires examining all possible prior probability beliefs and also all possible paths leading up to the prior probability belief. This process provides an extremely accurate characterization of uncertainty, however, for planning, the computation is intractable, as it is exponential in the size of the environment. We have therefore made some simplifying assumptions in the implementation of the algorithm which dramatically reduce the complexity. One such simplification allows us to ignore the problem of the prior positional distribution.

Figure 1 shows an example map on the left of the Smithsonian Museum of American History, and the information content of the same museum on the right. In the information content map, the darker an area is, the less information it contains. Notice that the darkest area is the center of the large open space in the middle, and that the lightest areas, with the lowest entropy are close to the walls.

Future Work: One avenue for future research lies in the path planner. The dynamic programming technique currently used for finding the minimum-cost trajectories demands a monotonic integration of the entropy. Therefore, there is no way to model actions that reduce uncertainty. Extending the planner to model uncertainty explicitly as a third state rather than including the information content in the cost function given to the planner would allow more formal modeling of the uncertainty and the application of this work to other domains. In particular, the current approach cannot be used for such traditional planning problems as mail delivery or scheduling.

Selected References:

[Kaelbling et al., 1998] Kaelbling, L. P., Littman, M. L., and Cassandra, A. R. (1998). Planning and acting in partially observable stochastic domains. *Artificial Intelligence*, 101.

[Koenig and Simmons, 1996] Koenig, S. and Simmons, R. (1996). The effect of representation and knowledge on goal-directed exploration with reinforcement learning algorithms. *Machine Learning Journal*, 22:227–250.

[Takeda et al., 1994] Takeda, H., Facchinetti, C., and Latombe, J.-C. (1994). Planning the motions of mobile robot in a sensory uncertainty field. *IEEE Trans. on Pattern Analysis and Machine Intelligence*, 16(10):1002–1017.