Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science

6.834J COGNITIVE ROBOTICS

Wednesday, May 12, 2004

Personal Insights in Course Project

By Morten Rufus Blas
rufus@mit.edu

Introduction:

The aim of this report is to describe the personal insights gained by the author in the final project of course 6.834J Cognitive Robotics. The projects goal was to create a SLAM (Simultaneous Localization and Mapping) tutorial for an indoor mobile robot. This tutorial is divided into a number of chapters pertaining to different problems in implementing a SLAM solution. As such the personal insights will be divided into a number of parts resembling the tutorials structure. These are: landmark extraction, data association, the EKF. I will then discuss briefly in more general terms the success of the actual tutorial in a conclusion.
Lastly this report is only meant read by someone who already has in-depth knowledge about SLAM.

Landmark Extraction:

Data collected from a laser scanner is very clean (accurate). As such one should think it trivial to implement a suitable landmark extraction policy to get good landmarks in an indoor environment. This proved not to be the case.

Some of the problems we encountered:

- Landmarks would be found close together and when the robot made for example a hard right turn, the left landmark would get associated to the right landmark (effectively killing our EKF as associations were wrong).

- We saw many papers extracting corners as landmarks but found it hard to find these in practice as corners look very different depending on the angle at which they are viewed. We initially were looking for corners formed by two walls meeting at for example 90° but often you would only be able to view one wall at a time. This made this type of landmarks characteristics from a laser scan hard to determine.

- People moving about in the environment could easily get flagged as landmarks. Our initial approach (Spikes) where we just looked at extremas was very bad at this as a person standing in front of a wall would be characterized as an extrema. Also how do you differ a person standing against a wall from a corner in the wall? If your landmark extraction algorithm looks for corners between small wall sections this can be very hard.

We decided to solve the majority of these problems by implementing a line extraction algorithm using RANSAC (Random Sampling Consensus). Strictly speaking RANSAC is just a method to determine the best fit line of a bunch of points (similar to least-squares). As such we needed to modify it to find a number of lines from a single laser scan. We have explained the algorithm in the paper but our innovation compared to what we have been able to find elsewhere are the following:

- Instead of sampling random points from the entire set of points to find one line you find local lines by randomly selecting one point from the whole laser scan sample and then getting your sample from points close to this point.

- Each time you successfully find a line you remove the points lying on the line from the set of points found by the laser scan. This prevents you from finding the same line twice and also decreases the size of the search space.

- We added a new stop criteria that says if there are only so many data points left from the laser scan then don't bother to try and extract lines.

RANSAC alleviated many of the problems in landmark extraction as:

- All buildings have walls (and most buildings have straight walls). Thus RANSAC is able to extract landmarks in all locations in a building.

- Walls are usually far apart so the problem of data-association is made much easier.

- It is quite robust towards people moving dynamically around in the environment. You can just specify the parameters in the RANSAC algorithm so it only picks up walls which are longer than the width of a human. In a very cluttered scene with many people this might not hold.

- RANSAC is quite robust towards occlusion so even if there is an object in front of the wall or people moving up and down partly obscuring the wall it can still be extracted.

However we also realized the following:

- The RANSAC algorithm is random. As such you might not always find all lines in a laser scan. If you are unlucky you may keep sampling the same points. In the implementation you get rid of this by resampling lots of times. This however is not very computationally efficient.

- The way we extract points on the RANSAC lines for use in the EKF is ok for our tutorial but has obvious limitations. By taking the point on the line closest to some fixed point like the origin you will unintentionally get most of these points to be close to each other (as they are all close to the origin). This makes data-association hard. Another solution may be to look at where the RANSAC lines intersect and use these as the fixed points. Both of these solutions however aim to boil down the information provided by a RANSAC line. The optimal solution would be to allow the EKF to use the whole line instead of just a point on the lines.

Data-association:

As previously described using RANSAC got rid of many of the data-association problems. We also described how to validate associations using the EKF innovation covariance.
Typical data-association strategies involve using the nearest-neighbor approach. Often the mahalanobis distance is used for this. As landmarks are far from each using our RANSAC method we found that using Euclidean distance also provides an adequate solution.

The EKF:

The Extended Kalman Filter in SLAM is described in many papers and was for our purposes well defined without much room for insights. The most taxing aspect was in figuring out how it actually works as what the different matrices in SLAM EKF actually contain was often times left out and assumed known a priori to reading a paper.

Insights which we gained in the EKF (beyond how actual updates are done) had most of all to do with the overall operation of the EKF. These aspects were either not documented or only briefly mentioned in the various SLAM papers that were read.

They are:

- As it does not matter whether you re-observe landmarks or add new landmarks first, you should always add new landmarks last as re-observing landmarks works on the whole matrix so is more computationally expensive if there are more landmarks in the covariance matrix.

- Adding errors in the EKF is important. Adding too little error will produce singular matrices which will cause the EKF to fail (as you cannot take their inverse). Learning to deal with this by adding bigger errors than what you would actually expect was part of the process of getting to know EKF and was an important insight.

- A minor insight was that it did not matter if you added new landmarks one at a time or you added all new landmarks at the same time. This facilitated our implementation as it was easier to see what was going on when you added a new landmark one at a time.

Conclusion:

The project as such was useful and provided valuable insight into the SLAM process. I learned a considerable amount. Before starting the project we did not have a very clear idea beyond what had been covered in class about the workings of SLAM. We did however realize the importance for a good tutorial in SLAM as we had read many papers on the subject prior to formulating our project description. Our main contribution was to remove much of the ambiguity present in many of the papers and especially in the EKF. I think this is a valuable contribution to the field as it should make the field accessible to people who are not SLAM-experts and people who don't have access to a SLAM-expert. My personal insights in landmark extraction and data-association have been valuable to me and have stimulated my interest to use cognitive robotics in a master thesis.