Kalman Filtering for Full Size Helicopter

Jeffrey Spehar

2007-12-14

CS229 Project Report

Project Description

This project is part of an autonomous flight project in the AI lab sponsored by Professor Ng. The current project has two main components. The first is to build fixed wing aircraft that are able to fly autonomously in large formations. The second involves autonomous flight control of a full sized helicopter. My work this quarter, under the supervision of Pieter Abbeel, focused on augmenting the code associated with localization and state estimation, integrating and testing the sensor systems with an existing Kalman Filter.

Sensor Integration and Testing

For the first several weeks, I worked on integrating sensor input and testing the Kalman Filter throughput. Both the full size helicopter and fixed wing aircraft will use the Kalman filter for localization and state estimation. Plans for the helicopter include several sensors for localization purposes. Among them are a GPS sensor and an IMU sensor. I took code that parses the raw input from a GPS device on a serial port and integrated it with the Kalman filter. The code for parsing the IMU sensor data was already written. I then tested that the Kalman Filter was working correctly with the physical device. I was able to confirm that the Kalman Filter works correctly with both the physical GPS and IMU sensing devices at a speed of approximately 150 state updates per second. This should be sufficiently fast to run the filter in real time while performing online localization.

Learning Filter Parameters with EM

With the sensor code in place and able to collect data, it was possible to move on to optimizing the parameters for the Kalman filter. The parameters were learned using the EM algorithm which, in the Kalman filter, can be expressed as running successive forward-backward passes that filter then smooth the sensor data. This produces a maximum likelihood estimation of the optimal filter parameters. In order to run EM, all the intermediate data about the state distributions from the forward filtering pass must be saved for use in the backward smoothing pass. For even small data sets, this data became too large to fit into memory (intermediate data from a 5 minute flight is nearly 1 GB). To solve this problem, I designed and implemented an efficient disk-based data structure to store the intermediate data.

Future Work: Beyond EM & Maximum Likelihood Estimation

The state predictions from the filtered and smoothed data incorporate all recorded measurements. Since only the filtered data can be used for state estimation, ideally, the filtered data would match the state predictions that have been filtered and smoothed. Currently, the EM formulation attempts to do this by maximizing likelihood, but perhaps an additional procedure could make improvements. The proposed future work involves minimizing the KL divergence between the filtered and smoothed state prediction distributions by a numerical gradient search, starting with the filter parameters obtained by running EM. It is unknown if this will improve performance.