CS 287, Fall 2009
Problem Set #3 : State estimation and SLAM

Deliverable: 1 page (max!) pdf write-up by Friday December 4, 23:59pm. This is meant to be a very short PS.

NOTE: Please refer to the class webpage for the homework policy.

1. For each of an extended Kalman filter and a particle filter describe an example setting in which they would be a suitable solution method and why they would be preferrable over the other one.

2. What does "closing the loop" refer to in SLAM and what is so special about closing the loop in SLAM?

3. Ridge regression solves problems of the form

$$\min_x \sum_{i=1}^{m} \left((a^{(i)})^\top x - b^{(i)} \right)^2 + \lambda x^\top x.$$

Describe how full ridge regression can be seen as computing the mean of x with a Gaussian prior and measurements of the form $b^{(i)} = (a^{(i)})^\top x +$ Gaussian noise. Make explicit how to set up the measurement and dynamics model to make a Kalman filter find the ridge regression solution in an incremental fashion.

4. The more particles in a particle filter, the more accurate the resulting estimate of the distribution. Describe what is particularly problematic about using a single particle. (Hint: consider the effect of the measurements on the estimated distribution.)

5. For landmark-based SLAM, describe how one might incorporate *not observing* a landmark when using a particle filter to represent the distribution over robot pose and landmark positions? Would a standard KF observation update be suitable in the EKF setting? Why (not)?