Albert-Ludwigs-Universität Freiburg
Lecture: Introduction to Mobile Robotics
Summer term 2007

Institut für Informatik

Prof. Dr. W. Burgard
MSc AI J. Sturm
Dipl.-Inf. B. Lau

Sheet 11

Topic: Monte Carlo Localization II

Submission deadline: Fri 13.07.2007, 11:00 a.m. (before class)

Introduction

In this exercise the landmark measurements are integrated into the particle filter using the likelihood function. It builds on exercise sheet 10, the solutions and the corresponding source code.

Each particle i with state `part(i,:)` is assigned a weight `weights(i)`. In case no resampling takes place, the weight w of one particle with state x is updated according to the following equation, where η is chosen such that $\sum_i \text{weights}(i) = 1$:

$$w_t = \eta \cdot w_{t-1} \cdot p(z_t|x_t), \tag{1}$$

Exercise 1:

In every time step a measurement for only one landmark is reported. It is always the one closest to the robot, however, this knowledge is not integrated into the likelihood function. Since the landmarks have no signature, the data association has to be estimated using the maximum likelihood selection. Implement the likelihood function based on the sensor model and the max-selection by completing the function `likelihood` in this week's stub. Implement Eq. 1 using `likelihood` in the main loop after the evolve step. Test it with uniform particle initialization. Use different numbers of particles to see how it influences the quality of the approximation.

Exercise 2:

(a) How do the computational costs of the particle filter scale depending on the number of particles and the number of dimensions in the state vector of the particles? Why can an increase of the dimensionality be a big problem in practice?

(b) How does resampling change the distribution of particles and the belief about the robot's position? How does resampling relate to the "kidnapped robot" problem, where the robot is placed at a different position withing notice, after localization has converged?