

LOCALIZATION WITH A PARTICLE FILTER


Tom Wilmots, John Larkin, and Ursula Monaghan

PARTICLE FILTER REVIEW

- Represent possible robot positions as (x,y,θ)
- Start with n particles randomly distributed in maze
- Motion Update:
 - Given control u , update particles: $P' = \{x_1, x_2, \dots, x_n\}$ with each x_i' sampled from $P(x_i' | x_i, u)$
- Measurement Update:
 - $W_i = \eta P(z | x_i')$
- W_i is the probability that particle i is a good particle

RAY TRACER

- Built a Ray Tracer to determine our expected measurements
- Ray Tracer takes a map of the world and sends out rays from a point and tests intersections with the world
- We record the length of the ray that hits the closest object



$$P_i = P_o + \alpha d$$

$$P_i = P_1 + \beta(P_2 - P_1) = P_1 + \beta k$$

where $k = (P_2 - P_1)$



$$P_1 + \beta k = P_o + \alpha d$$

$$\alpha d - \beta k = P_1 - P_o$$

$$\begin{bmatrix} d & k \end{bmatrix} \begin{bmatrix} \alpha \\ \beta \end{bmatrix} = \begin{bmatrix} P_1 - P_o \end{bmatrix}$$

Solve for α to get distances, check to see if β is between (0,1)

VISUALIZATION OF RAY TRACER



REAL MEASUREMENT

- Subscribed to the SCAN topic of ROS to collect real measurements from Kinect
- Kinect was able to scan range of $\pm 30^\circ$ and output 640 rays, giving 640 distance measurements
- Did not use all 640 measurements, selected a subset of the rays to minimize computational runtime.
- Kinect is not at center of robot: Need to transform to account for the position of the scan

UPDATING WEIGHTS

Particles = 3 by n array (n particles with x,y, θ)

weights = list of size n

z = measured rays from kinect

for i in particles:

 list_of_rays = expected measurements for this particle from Raytracer

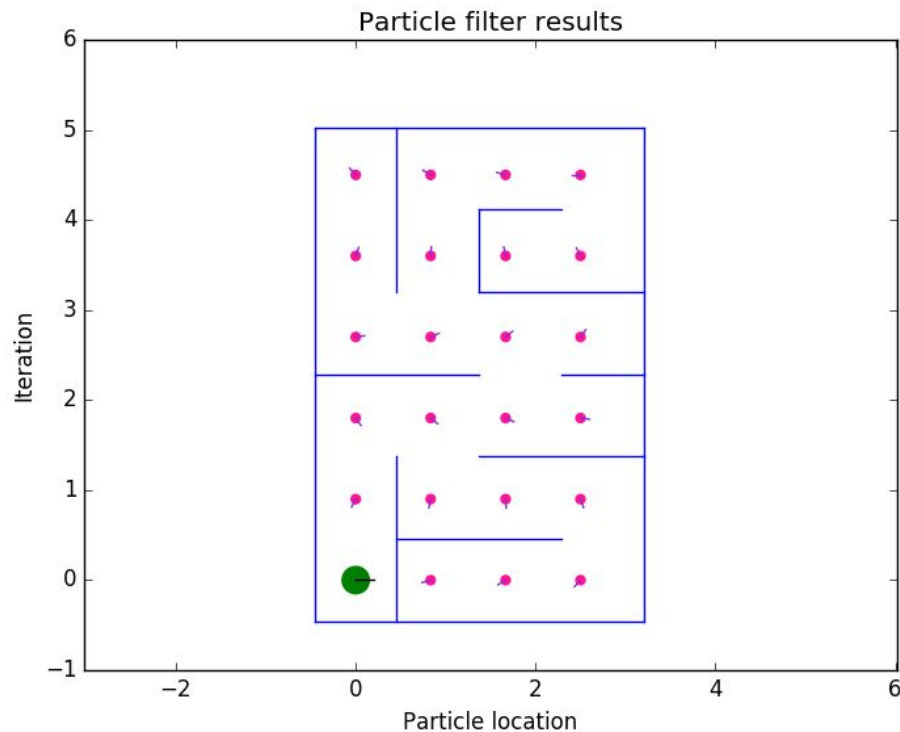
 for j in list_of_rays:

 weight[i] = weight[i] * $\exp((z[j] - \text{list_of_rays}[j])^2 / (2 * \sigma^2))$

 weight[i] = weight[i] / sum(weight[i]) #normalize weights

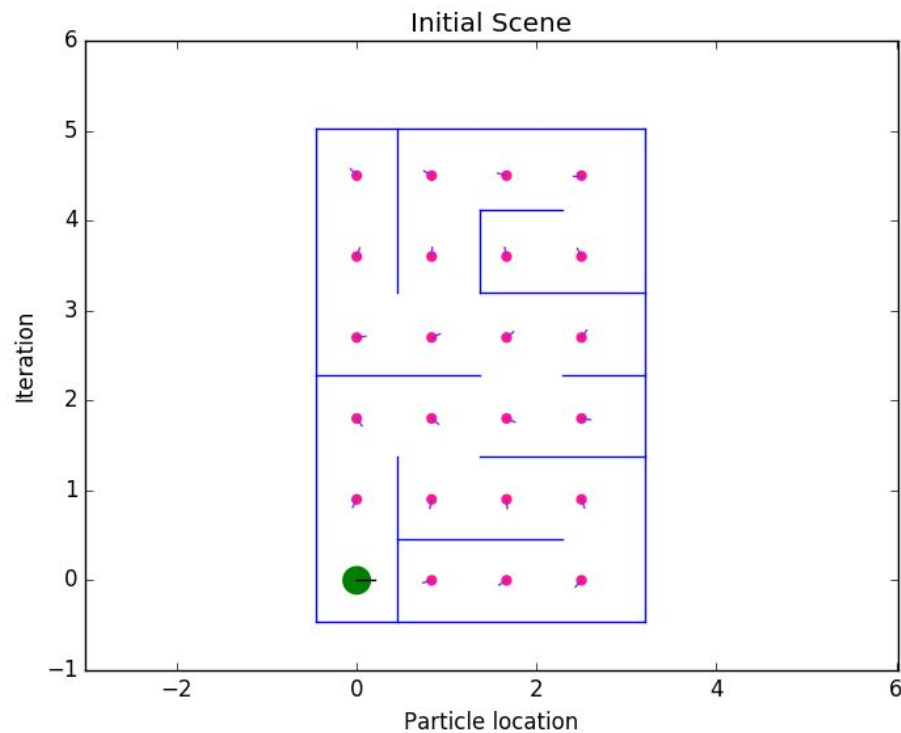
#resample particles with replacement

RESULTS (TRIAL 1)



- Measurement and motion step are plotted sequentially for a given iteration
- Pretty accurate path
 - During the run, Leela *actually* veered near cell (2,1)

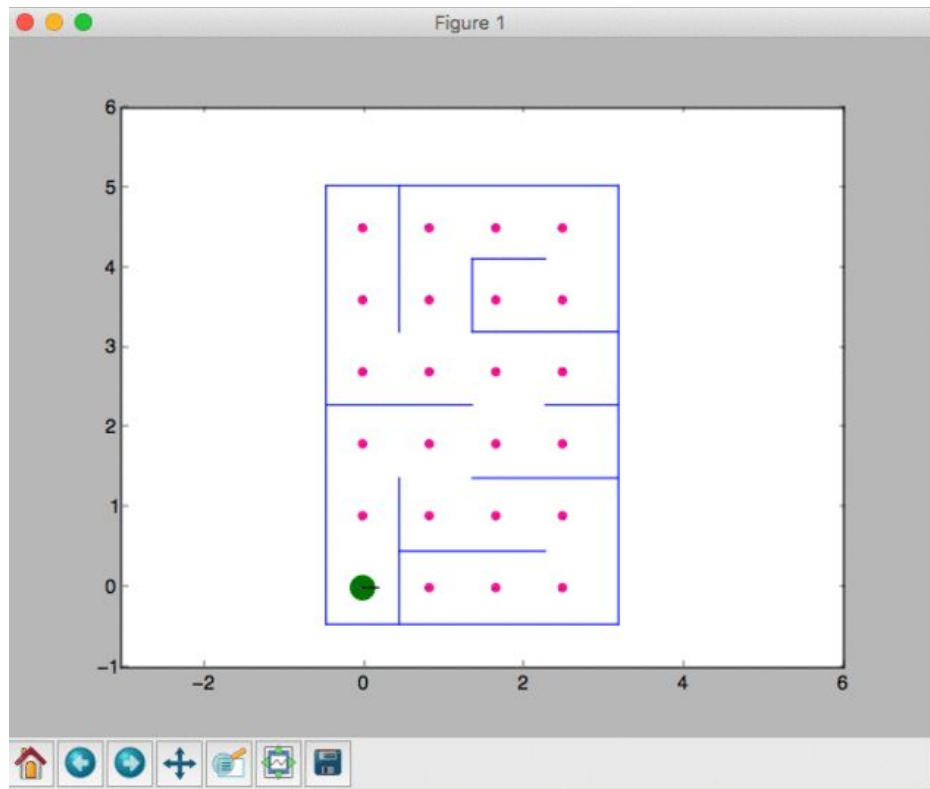
RESULTS (TRIAL 2)



- Same path as before

PROBLEMS

- Ran into a lot of different bugs
 - Motion Update
 - Underflow
 - Additive noise with motion update
 - Resampling
 - Real potential for non-convergence
 - Correct computation of weights /
fundamental understanding of particle filter



EXTENSIONS

- Optimizing runtime - believe to be $O(n*m*w)$ where n is the number of particles and m is the number of rays and w is the number of walls
 - More advanced data structures could reduce runtime
 - Could allow for more particles
- Degree of scan
 - Localization would be expedited with greater range of scan
 - Easy for raytracer
 - Harder for Leela