CS354

Nathan Sprague

September 28, 2023

Monte-Carlo Localization aka Particle Filter

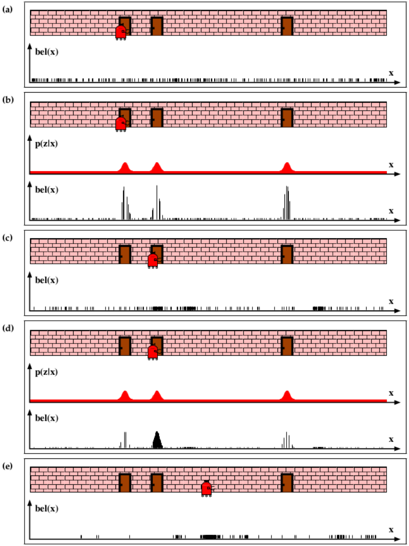

Figure 8.11 Monte Carlo Localization, a particle filter applied to mobile robot localization.

Probabilistic Robotics. Thrun, Burgard, Fox, 2005

Particle Filter Algorithm

procedure PARTICLE_FILTER_FOR_LOCALIZATION$(\mathcal{X}_{t-1}, u_t, z_t, m)$

  Inputs

    $\mathcal{X}_{t-1}$ – The previous set of particles

    u_t – The control signal

    z_t – The sensor value

    m – The map

  Output

    $\mathcal{X}_t$ – The updated set of particles

  for $m = 0$ to $M - 1$ **do**

    $x_t^{[m]} = \text{sample\_motion\_model}(u_t, x_{t-1}^{[m]})$ ▷ Predict

    $w_t^{[m]} = \text{measurement\_model}(z_t, x_t^{[m]}, m)$ ▷ Correct

    $\bar{\mathcal{X}}_t = \bar{\mathcal{X}}_t \cup \{\langle x_t^{[m]}, w_t^{[m]} \rangle\}$

  for $m = 0$ to $M - 1$ **do** ▷ Resampling

    draw i with probability $\propto w_t^{[i]}$

    $\mathcal{X}_t = \mathcal{X}_t \cup \{\langle x_t^{[i]}, 1/M \rangle\}$

Based on Algorithm in Table 8.2 in Probabilistic Robotics. Thrun, Burgard, Fox, 2005

Sampling From the Motion Model

- $x_t^{[m]} \sim p(x_t \mid u_t, x_{t-1}^{[m]})$

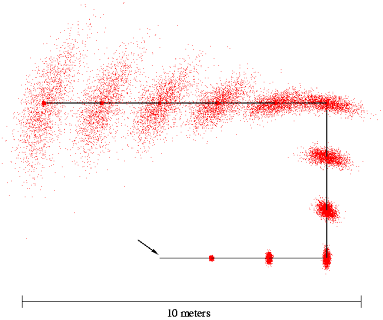

http://robots.stanford.edu/probabilistic-robotics/

Measurement Models for Laser Range Finders

- $w_t^{[m]} = p(z_t \mid x_t^{[m]}, m)$
- (Note that weights won't sum to one)

measurement noise:

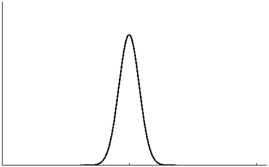

unexpected obstacles:

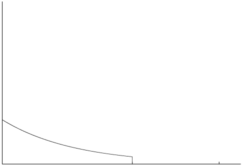

random measurement:

maximum range:

http://robots.stanford.edu/probabilistic-robotics/

Measurement Models for Laser Range Finders

- $w_t^{[m]} = p(z_t \mid x_t^{[m]}, m)$

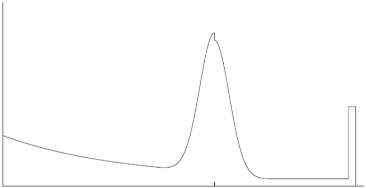

$$P(z \mid x, m) = \begin{pmatrix} \alpha_{\text{hit}} \\ \alpha_{\text{unexp}} \\ \alpha_{\text{max}} \\ \alpha_{\text{rand}} \end{pmatrix}^T \cdot \begin{pmatrix} P_{\text{hit}}(z \mid x, m) \\ P_{\text{unexp}}(z \mid x, m) \\ P_{\text{max}}(z \mid x, m) \\ P_{\text{rand}}(z \mid x, m) \end{pmatrix}$$

http://robots.stanford.edu/probabilistic-robotics/

Issues

- Not a good idea to run the particle filter while the robot is stationary
  - Resampling will deplete the set of particles
- May not be necessary to resample on every update.

Extracting a Single State Estimate

- Possibilities:
  - Average over all particles
  - Cluster the algorithms, average within the “best” cluster
  - Something fancier...